

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
20 February 2003 (20.02.2003)

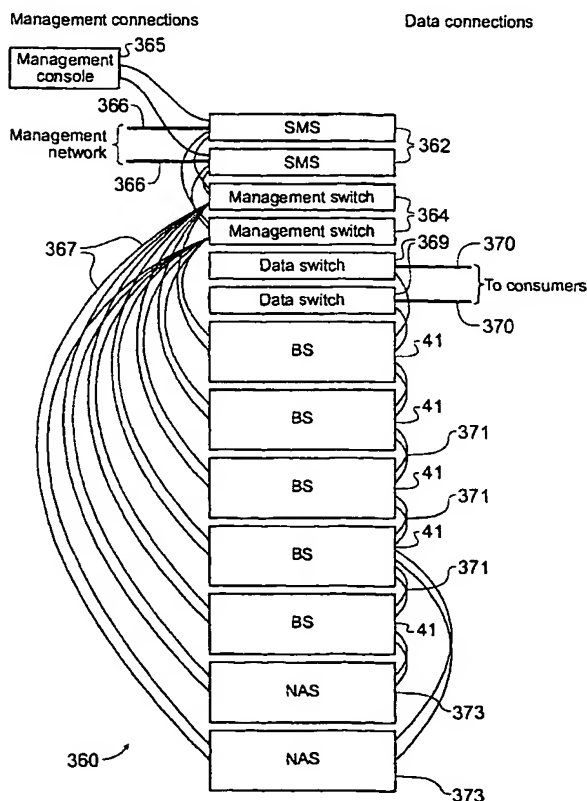
PCT

(10) International Publication Number
WO 03/014892 A2

- (51) International Patent Classification⁷: **G06F 1/16**
- (21) International Application Number: **PCT/GB02/03665**
- (22) International Filing Date: **9 August 2002 (09.08.2002)**
- (25) Filing Language: **English**
- (26) Publication Language: **English**
- (30) Priority Data:
- | | | |
|------------|-----------------------------|----|
| 60/311,505 | 10 August 2001 (10.08.2001) | US |
| 0119558.5 | 10 August 2001 (10.08.2001) | GB |
| 60/366,430 | 21 March 2002 (21.03.2002) | US |
| 0206719.7 | 21 March 2002 (21.03.2002) | GB |
| 10/172,722 | 14 June 2002 (14.06.2002) | US |
- (71) Applicant (for all designated States except US): **SUN MICROSYSTEMS, INC** [US/US]; 901 San Antonio Road, Palo Alto, CA 94303 (US).
- (72) Applicants and Inventors (for US only): **KING, James, E.** [GB/GB]; 75A Waterloo Road, Wokingham, Berkshire RG40 2JE (GB). **MAYHEAD, Martin, P.** [GB/GB]; Tuscaloosa, Tilford Road, Hindhead, Surrey GU26 6QY (GB).
- (72) Inventor; and (75) Inventor/Applicant (for US only): **GARNETT, Paul, J.** [GB/GB]; Little Orchard, Knightsbridge Road, Camberley, Surrey GU15 3TS (GB).
- (74) Agents: **HARRIS, Ian, Richard** et al.; D Young & Co., 21 New Fetter Lane, London EC4A 1DA (GB).
- (81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG,

[Continued on next page]

(54) Title: **SYSTEM MANAGEMENT**



(57) Abstract: A server system comprises a server blade carrier. The server blade carrier carries a plurality of removably mountable server blades. Each blade includes a blade service controller operable to perform monitoring and management functions in the server blade. The carrier comprises at least one carrier service processor. The carrier service processor is operably connected to the blade service controllers of blade servers mounted in the carrier to provide higher level management functions and is configured to communicate management functions to and from the blades.

WO 03/014892 A2



SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ,
VC, VN, YU, ZA, ZM, ZW.

(84) **Designated States (regional):** ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

SYSTEM MANAGEMENT

BACKGROUND OF THE INVENTION

5 The present invention relates to system management, and in particular but not exclusively to management of multi-processor systems, for example multi-processor server systems.

One application for the present invention relates to high density computer
10 systems, for example, computer server systems for telecommunications applications. In telecommunications applications, it is important to provide high reliability and high capacity of operation. Various approaches have been taken to providing such high-performance, high reliability systems. Typically such systems are designed around providing redundant resources so that if one component of the system develops a
15 fault, the system remains operational using the redundant resources. Fault tolerance can also be achieved, for example, with multiprocessor systems that provide redundancy through dynamic, e.g., software-controlled, task distribution. High density systems are typically rack mountable, with one or more processor systems occupying a shelf in the rack. The trend in recent times is to make the computers with
20 smaller form factors. This means that more computers can be located in a rack. This has the advantage of increasing the processing density within the racks, and also the advantage of reducing the distance between the computer systems.

The present invention relates to providing system management functions for a
25 modular computer system.

SUMMARY OF THE INVENTION

A first aspect of the present invention can provide a server system that includes a server blade carrier. The blade carrier can carry a plurality of removably mountable blades. Each blade can include a blade service controller operable to perform monitoring and management functions in the blade. The carrier can comprise at least one carrier service processor. The carrier service processor can be operably connected to the blade service controllers of blade servers mounted in the carrier to provide higher level management functions and can be configured to communicate management functions to and from the blades.

An embodiment of the invention can provide a hierarchical distribution of management functions within the server system, allowing management functions to be distributed to an appropriate level within the system. This management structure is readily scalable as the system grows. Also, it provides cost effective system management. The lower level management functions provided at a lower level have lower resource requirements to perform them than higher level functions provided at a higher level. It further enhances field replacement of blades and the insertion, in use, of new blades as a blade to be inserted include a service controller.

20

One embodiment provides the server system additionally comprising a system management server. The system management server can be operably connected to the carrier service processor of at least one blade carrier and the blade servers mounted in the carriers to provide high level management functions. This arrangement provides an additional tier to the management hierarchy allowing further distribution of management functions.

A further aspect of the present invention provides a method of providing management functions within a server system having a server blade carrier carrying a

plurality of removably mountable server blades, the method comprising: performing monitoring and management functions at a blade service controller within each blade; performing higher level management functions at at least one carrier service processor within the carrier; providing a communications channel connecting the carrier service
5 processor to each blade service controller; and communicating management functions between the carrier service processor and the blade service controllers.

Further aspects and advantages of the invention will become apparent from the following description of particular embodiments.

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments of the present invention will be described hereinafter, by way of example only, with reference to the accompanying drawings in which like reference
5 signs relate to like elements and in which:

Figure 1 is a schematic representation of an architecture of a multiprocessor system for supporting a web site;

10 Figure 2 is a schematic representation of a racking system incorporating an example of a carrier in the form of a rack-mountable shelf according to a first example;

Figure 3 is a front view of an example of a carrier in the form of a shelf of
15 Figure 2;

Figure 4 is a rear view of an example of the shelf of Figure 2;

Figures 5A, 5B and 5C are schematic perspective views and a schematic
20 exploded view respectively of an example of an information processing cartridge for mounting in the shelf of Figure 2;

Figure 5D is a schematic perspective view of an example of an information processing cartridge for mounting in the shelf of Figure 2;
25

Figure 6 is a schematic perspective view of an example of a combined switch and service processor module for mounting in the shelf of Figure 2;

Figure 7 is a schematic perspective view of an example of a power supply module for mounting in the shelf of Figure 2;

Figures 8A, 8B and 8C are a schematic plan view and schematic perspective views, respectively, of an example of the chassis and midplane of the shelf of Figure 2;

Figure 8D is a schematic perspective view of a part of a base member of the shelf of Figure 2;

Figures 9A, 9B and 9C are schematic front, top and rear views, respectively, of an example of a midplane of the shelf of Figure 2;

Figure 10 is a schematic cross section view through the shelf of Figure 2;

Figure 11 is a functional block diagram of an example of an information processing subsystem for the information processing cartridge of Figure 5;

Figure 12 is a functional block diagram of an example of an information processing subsystem for the combined switch and service processor module of Figure 6;

Figure 13 is a functional block diagram of an example of a subsystem for the power supply unit of Figure 7;

Figure 14 is a functional block diagram showing the connectivity between the components of the shelf of Figure 2;

Figure 15 is a functional block diagram showing the external connectivity of the shelf of Figure 2;

Figure 16 is a schematic representation of a shelf showing the external
5 connections from the shelf of Figure 2;

Figure 17 is a schematic representation of a rack mounted system comprising a plurality of such shelves; and

10 Figure 18 is a functional block diagram of a further example of an information processing subsystem for the information processing cartridge of Figure 5.

While the invention is susceptible to various modifications and alternative forms, specific embodiments are shown by way of example in the drawings and are
15 herein described in detail. It should be understood, however, that drawings and detailed description thereto are not intended to limit the invention to the particular form disclosed, but on the contrary, the invention is to cover all modifications, equivalents and alternatives falling within the spirit and scope of the present invention as defined by the appended claims.

DESCRIPTION OF PARTICULAR EMBODIMENTS

Embodiments and examples are described hereafter by way of example only in the following with reference to the accompanying drawings.

5

Shown in Figure 1 is an example of an application of a high capacity multiserver system 1 for implementing a network-connected web site such as, for example, an airline reservation system on the World Wide Web.

10 As shown in Figure 1, an external network 3 (e.g., the Internet) for communicating with a user 5 can be connected to gateways 7 which can be connected to an entry edge server group 9 implemented by a web farm as shown in Figure 1. The entry edge server group 9 forms an interface to the external network 3. The entry edge server group 9 can then be connected by switches 11 and a firewall 13 to a web
15 edge server group 15 that can also be implemented as a web farm as shown in Figure 1. The web edge server group 15 can serve to cache web pages that are readily accessible to users 5 accessing the system 1 from the external network 3, for example for checking flight times, etc. The web edge server group can comprise a number of blade server (BS) shelves and a number of network addressable storage (NAS) shelves
20 for storing critical data. The web edge server group 15 can be further connected by a further firewall 17 to a plurality of application servers 19, which can be responsible for, for example, processing flight reservations. The application servers 19 can then be connected via a further firewall 21 to computer systems 23, 25, for example, e-commerce services including financial services for receiving and processing payment
25 for airline reservations.

As will be appreciated, the server system described above with reference to Figure 1 is only an example of a possible application for a multiprocessor server system. Multiprocessor server systems have many different applications and the

present system is not limited to being applicable for use in only one or a limited number of such applications, rather multiprocessor server systems as described herein are operable for use in many different applications. A non-exhaustive list of such alternative applications includes: e-commerce web server systems; 5 telecommunications network server systems; LAN application and file server systems and remote vehicle control systems.

With reference to Figure 2, there is shown a schematic perspective representation of a rack system 31 as viewed from the front including left and right 10 front uprights 32 and 33 and left and right rear uprights 34 and 35. The uprights can be formed with apertures for receiving shelf fixings (e.g., screws, bolts, clips, etc., for mounting brackets, slides, rails, etc.).

Also shown in Figure 2 is an example of a blade server shelf 41 mounted in 15 the rack system 31. The shelf 41 forms a carrier configured to carry a plurality of information processing cartridges 43 located side by side along the shelf.

The term "shelf" is used herein in a conventional way to describe a structure that is mountable in rack system 31 and is configured to carry one or more 20 components to form at least a part of a rack-mountable system. In the present example, the shelf 41 is three-dimensional, having a height (H), width (W) and depth (D). In the present example, one dimension (hereinafter described as the height, H) is smaller than the other dimensions (hereinafter described as the depth, D, and the width, W) to facilitate mounting of the shelf within the rack system 31. It will be 25 appreciated that although the width and depth are typically constrained by the dimensions of the racking system for which the shelf is designed, there is more freedom as regard the height, subject to taking account of any appropriate standards and packaging considerations.

Each of the information processing cartridges contains at least one processor. Each information processing cartridge in the present example is operable as a server. In the described examples, the information processing cartridges are configured as robust enclosed modules.

5

In the example to be described in more detail presently, the information processing cartridges, when aligned in the carrier shelf, look like rectangular slabs, or blades. Accordingly, an information processing cartridge can be described as a blade.

The information processing cartridges 43 comprise information processing modules
10 enclosed in an enclosure, or housing, so that the information processing modules have the form of cartridges. Also, as the information processing cartridges are to operate as computer servers in the example described in more detail presently, an information processing cartridge 43 can also be described as a server blade. Accordingly, in the context of this example, the terms module, cartridge and blade are used
15 interchangeably.

The illustrated example of a shelf 41 is configured to carry sixteen information processing cartridges 43, each of which is removably mountable in a respective opening 45 in the front of the shelf, whereby the information processing cartridges can
20 be inserted into and removed from the front of the shelf 41 without removing the shelf 41 from the rack system 31.

In the present example, the shelf 41 comprises a three-dimensional, generally rectangular, enclosure, or housing, 47 that is suitable for mounting in generic racking
25 systems including both 4-post and 2-post systems. It can be mounted on fixed rigid rack mounting ears and/or a simple slide/support system. The present example is designed for standard 19"-wide racking (1"=25.4mm) as defined, for example, in the well-known IEC297 and EIA 310 specification standards with height corresponding to the so-called 3U (3 standard unit) height. For mounting such a 3U unit in such a 19"-

wide racking system, with a depth of, say 25" or 30", the enclosure can be arranged with a height of up to about 130.5mm, a width of up to about 445 mm and a depth, including all hardware and fascias, but excluding cable management, of up to about 635mm, with the depth from the front-most point of a fascia to a rear I/O connector panel of a rear mounted Field Replaceable Unit (FRU) of about 610mm. Of course, other examples designed for other racking systems could have different dimensions.

This example of a shelf 41 has a single enclosure, or housing, 47 that houses a number of modular units or subsystems, the majority of which are replaceable in the field and are therefore known as Field Replaceable Units (FRUs). These modular units include the information processing cartridges 43.

The shelf enclosure 47 can be fabricated from sheet material (e.g., from steel sheet) to form a chassis portion 49 that includes a base 51, two sides 53 and 55, a front 57 and a rear 59. The word "front" as used here is merely used as a label herein to refer to the face, or wall 57 of the enclosure that is located at the main access side of the rack system 31 in use when the shelf is mounted therein. Similarly, the words "rear" and "side" are merely used as labels herein to refer to the faces, or walls 59, 53 and 55 that, in use, are located at those respective positions when the shelf is mounted in the rack system 31.

The openings 45 can be formed in the front face 57 for receiving the information processing cartridges 43 and, as will be explained later, apertures can also be formed in the rear face 59 for receiving further FRUs. The enclosure can further include a removable top cover 61 that can be secured to the chassis portion 49 by suitable fastening (e.g., screws). The apertures in the front and rear faces 57 and 59 allow at least some of the FRUs to be inserted into and/or removed from the shelf enclosure 47 via the front or the rear thereof, as appropriate, without removing the shelf from the racking. Access to components mounted in the shelf that are not

accessible via one of the apertures in the front 47 and rear 59 faces can be achieved by removing the shelf enclosure 47 from the racking system 31 and then removing the top cover 61 of the shelf enclosure 47.

5 Figure 3 is a front view of an example of a shelf 41 for a first example. A plastic front bezel 63 can be provided that fits on the front face 57 (shown in Figure 2) of the chassis 49 of the shelf enclosure 47. The front bezel 63 can be formed as a unitary removable part that spans the whole width and height of the front of the shelf enclosure 47. The front bezel 63 could alternatively include a number of separate
10 components, or mouldings. The front bezel can include a peripheral portion 64 that can provide areas for corporate and product branding marks, for identification and numbering for the information processing cartridge and for a bar code label (all not shown). One or more apertures 65 can be formed in the peripheral portion 64 of the bezel 63. The apertures 65 in the bezel can be arranged to align with one or more
15 apertures (e.g. a slot (not shown in Figure 3) in the front face of the chassis. In use, air can pass through the apertures 65 to flow into the shelf enclosure 47 to reach FRUs that are mounted in the shelf enclosure 47 through the rear face 59 thereof. Air flowing through the aperture 65 flows into a plenum chamber 66 (not shown in Figure 3) to flow past the processing cartridges 43 to reach rear mounted FRUs. A central
20 area 67 of the front bezel 63 can be open allowing access to the openings 45 in the front face 57 of the shelf enclosure 47 for insertion and removal of the information processing cartridges 43. Where no active module is mounted in a location for an information processing module, a blanking panel, or filler panel, such as the blanking panel 44, can be located in that location. LED indicators 69 can be mounted on a
25 system indicator printed circuit board (not shown) behind a designated area of the bezel to provide an indication of system status via light guides incorporated into the bezel. A further system indicator board (also not shown) carrying LED indicators can be provided inside the shelf enclosure to be visible from the rear thereof.

As mentioned above, in the present example of a shelf, up to sixteen information processing cartridges 43 can be installed in respective openings 45 in the front face 57 thereof. The number of information processing cartridges 43 actually installed in any installation is dependent upon the system configuration required.

5 Various features relating to the information processing cartridges 43 that are shown in Figure 3 will be described later.

Figure 4 illustrates the rear of the shelf unit of Figures 2 and 3. This shows two different types of FRU 71 and 81 (4 units in total) that have been inserted into

10 respective apertures 72 and 82 in the rear of the shelf enclosure 47. The FRUs shown in Figure 4 include two Combined Switch and Service Processors (CSSPs) 71 and two Power Supply Units (PSUs) 81. Various features shown in Figure 4 will be described later.

15 Before proceeding with a more detailed description of each of the FRUs 43, 71 and 81 introduced so far and of the construction of the shelf 41, there follows a brief description of an information processing cartridge 43, a CSSP 71 and a PSU 81 with reference to Figures 3, 4, 5, 6 and 7.

20 Figure 5A provides a perspective view, partly from the rear, of an information processing cartridge 43. Figure 5B provides a perspective view, partly from the front of the same information processing cartridge 43. Figure 5C provides an exploded perspective view of the construction of the information processing cartridge 43. Here is it to be noted that the term "rear" is applied in the context of the position, when

25 installed, of the information processing cartridge, with respect to the shelf 41 (i.e. in this case the "rear" of the information processing cartridge 43 is the innermost part of the information processing cartridge when it is inserted in the shelf 41). Likewise, "front" refers in the present context to the outermost part of the information processing cartridge when it is inserted in the shelf 41.

With reference to Figures 3, 5A and 5B, it will be noted that information processing cartridges are three-dimensional, having a height (h), width (w) and depth (d). If, as in the present example, the information processing cartridges are to be
5 arranged in a one-dimensional array (a row) across the shelf, then efficient packing for the information processing cartridges is achieved where one dimension (here the width, w) is smaller than the other dimensions (here the depth, d, and the height, h). In a particular example, the enclosure of a processing cartridge 43 has a height h, width w and depth d of 115mm, 26mm and 315mm, respectively, although of course,
10 other examples may have different dimensions.

It will be noted that an enclosure 101 of the present example of an information processing cartridge 43 has six, generally rectangular, faces. For convenience only, the face that is visible from the front of the racking when an information processing
15 cartridge 43 is mounted in the shelf 41 is known as the front face 102. The opposite face is known as the rear face 103. In the present example these two faces, as well as top and bottom faces 104 and 105, and the side faces 106 and 107 have the shape of elongate rectangles.

20 Although in the present example the information processing cartridges have six generally rectangular faces, it will be appreciated that other examples could have other configurations. For example, rather than having a generally rectangular shape, the side faces of another example of an information processing cartridge could have the general shape of a triangle (whereby the information processing cartridge may then
25 only have five faces), a pentagon (whereby the information processing cartridge may then have seven faces), and so on. Indeed, one or more or all of the edges could be curved. However, it will be appreciated that the present configuration provides advantages for example, in terms of manufacturing, engineering and packing density within a shelf 41.

In this example, the information processing cartridge enclosure 101 is fabricated from two housing portions. The first housing portion 109 can be fabricated from, for example, a metal (e.g., pressed steel) and can include one side face 107, the rear face 103 and part of the top face 104. The part of the top face formed by the metal portion 109 is given the reference 1041 in Figure 5A. The second housing portion 110 can be fabricated from, for example, a plastics material and can include the other side face 108 and the front faces 102 and bottom face 105 and the remaining part 1042 of the top face 104. In the present example, a plastics material used is PolyCarbonate Acrylonitrile Butadiene Styrene (PCABS), however many other plastics materials such as other ABSs and nylons may be used. The metal portion 109 also includes a part 1021 corresponding to the front face 102 and a part 1051 corresponding to a part of the bottom face 105. The parts are covered by the plastics portion 110 when enclosure 101 is fully assembled. These parts are visible in the exploded view of Figure 5C. The two housing portions 109 and 110 are secured to one another by fixings such as screws 118.

As shown in Figure 5A, a groove 108 runs along the length of the plastic top face portion 1042. This groove 108 is provided for interfacing with a guide member of the shelf chassis (not shown in Figure 5A, but see Figures 8B and 8C). A similar groove 113 is formed in the bottom face 105 of the plastics portion 110 as shown in Figures 5B and 5C. This groove 113 is provided for interfacing with a guide member of the shelf chassis (not shown in Figures 5B and 5C, but see Figure 8D).

The provision of the enclosure 101 means that the information processing cartridge 43 can safely be handled by an operator who is not a skilled technician. Also, through the use of the enclosure 101, the information processing cartridge is a robust unit that protects its inner workings from the outside environment and vice versa. The use of a conductive enclosure, e.g., a metal enclosure, can result in the

information processing cartridge including its own electromagnetic shielding. To this end, the enclosure 101 can also be provided with an internal electromagnetic interference (EMI) shielding member 116 as shown in Figure 5C. The EMI shielding member can be located between the internal components 112 of the information
5 processing cartridge 43 (not described in detail with reference to Figures 5A to 5C, but see Figures 11 and 18 below) and the plastics portion 110. The EMI shielding member can be secured to the plastics portion 110, for example by heat bonding or an adhesive. In other examples, the plastics portion could have a conductive layer deposited thereon, or conductive plastics material could be used. In this example, the
10 EMI shielding member 116 is provided with electromagnetic interference (EMI) fingers 114 to ensure good contact with the shelf chassis and the adjacent components. These fingers 114 extend through EMI finger apertures 1141 in the plastics portion 110.

15 As shown in Figures 3, 5B and 5C, the information processing cartridge 43 incorporates an injector/ejector handle 111 on the front face 102 for facilitating insertion and latching of the information processing cartridge 43 within an aperture in the shelf 41.

20 As shown in those Figures, the handle 111 of the injector/ejector lever extends substantially the whole height of the front face of the information processing cartridge 43, thereby increasing the mechanical advantage and facilitating injection and ejection of the information processing cartridge 43. As is further shown in those Figures, the front face 102 of the information processing cartridge 43 has perforations 115, in the
25 present example slits, to allow for airflow into the information processing cartridge 43. The front face part 1021 of the metal portion 109 has perforations corresponding to those in the front face 102 such that airflow into the information processing cartridge 43 is not impeded. It will be noted in Figures 3, 5B and 5C that the handle 111 is narrower in its middle than at its ends. This reduces any masking effect of the

handle 111 on the airflow to the perforations 115 in the front face of the information processing cartridge 43 and facilitates grasping of handle 111. The handle 111 can be bowed out from the front face of the information processing cartridge to further facilitate grasping thereof and to reduce any masking effect with regard to the
5 perforations. In other examples, the handle 111 could have an open frame-like structure to further facilitate airflow.

As shown in Figure 5A, the rear face 103 of the information processing cartridge 43 also has perforations 117 to allow for air to be exhausted from the rear of
10 the information processing cartridge 43. A fan can be located within the enclosure 101 of an information processing cartridge 43. In the present example the fan is a combined radial fan (blower) and heat sink to direct cooling air onto a processor of the information processing cartridge, which fan also encourages air movement within the enclosure 101. LED indicators 119 (see Figure 3 and 5B) can be provided on the front
15 face 102 of an information processing cartridge to indicate whether power is on, whether service intervention is required and whether the information processing cartridge 43 can be removed. A connector 120, for example a 40 way single connector attachment (SCA-2) connector (a small computer systems interface (SCSI) connector), can be provided at the rear of the information processing cartridge 43 for
20 electrical connection of the information processing cartridge 43 within the shelf 41. The connector 120 is advantageously able to withstand repeated removals and insertions of the enclosure 101 from and into a shelf 41. The connector arrangement can include a guide pin arrangement to prevent module misalignment during insertion of the information processing cartridge into the receiving location.

25

Thus an example of features and the construction of an information processing cartridge enclosure has been described. Although particular materials and constructions have been described, it will be appreciated that other examples could be employed. Indeed, it will be appreciated that this particular example relates to only a

possible form for the processing cartridge 43. An example of an alternative construction will now be described with reference to Figure 5D.

Figure 5D provides a perspective view, partly from the rear, of an alternative
5 information processing cartridge 43. Here it is to be noted that the term "rear" is applied in the context of the position, when installed, of the information processing cartridge, with respect to the shelf 41 (i.e. in this case the "rear" of the information processing cartridge 43 is the innermost part of the information processing cartridge when it is inserted in the shelf 41).

10

In this example, the information processing cartridge enclosure 101 is fabricated from pressed steel to form two chassis portions. The first portion 234 includes one side face 107, and part of each of the front and rear faces 102 and 103 and the top and bottom faces 104 and 105. The second portion 235 includes the other
15 side face 108 and the remaining part of each of the front and rear faces 102 and 103 and the top and bottom faces 104 and 105. The two chassis portions 234 and 235 meet at a groove 236 and are secured to one another by fixings (e.g., one or more screws, not shown). Grooves 236 run along the top and bottom faces 104 and 105 of the enclosure 101 and are provided for interfacing with guide rails of the shelf chassis
20 49 (not shown in Figure 5, but see Figure 8C). A cover portion that is secured to the chassis portion forms the other side face 106. It will be appreciated however, that in another example, the chassis portions could be joined at a position other than the groove 108, with the groove 108 being formed entirely in one of the chassis portions. Alternatively, the enclosure 101 may be constructed from a number of sheets of steel,
25 with each sheet forming one of the faces.

Figure 6 provides a perspective view, partly from the front, of a Combined Switch and Service Processor (CSSP) cartridge (or CSSP) 71. Here it is to be noted that the term "front" is applied in the context of the position, when installed, of the

CSSP cartridge 71, with respect to the shelf 41 (i.e. in this case the "front" of the CSSP cartridge 71 is the innermost part of the CSSP cartridge 71 when it is inserted in the shelf 41).

5 With reference to Figures 4 and 6, it will be noted that a CSSP cartridge 71 is three-dimensional, having a height (h), width (w) and depth (d). In a particular example, the enclosure of a CSSP 71 has a height h, width w and depth d of 43mm, 202mm and 278mm, respectively, although of course, other examples may have different dimensions.

10

 An enclosure 121 of present example of a CSSP cartridge 71 has six, generally rectangular, faces. For convenience only, the face that is visible from the rear of the racking when a CSSP cartridge 71 is mounted in the shelf 41 is known as the rear face 122. The opposite face is known as the front face 123. In the present example these
15 two faces, as well as side faces 126 and 127 have the shape of elongate rectangles. The top and bottom faces 124 and 125 are also rectangular, but not elongate in the manner of the front, rear, top and bottom faces. Although in the present example the CSSP cartridges have six generally rectangular faces, as for the information processing cartridges 43 it will be appreciated that other examples could have other
20 configurations.

 In this example, the CSSP enclosure 121 is fabricated from steel sheet to form a chassis portion that includes the bottom face 125, the front and rear faces 122 and 123 and the side faces 126 and 127. A cover portion that is secured to the chassis
25 portion forms the other top face 124. The cover portion is secured to the chassis portion by suitable fixings, for example one or more screws 128. It will be appreciated however, that in another example, other faces, or portions, of the enclosure could form the chassis and the cover portions. The provision of the enclosure 121 means that the CSSP cartridge 71 can safely be handled by an operator

who is not a skilled technician. Also, through the use of the enclosure 121, the switch cartridge is a robust unit that protects its inner workings from the outside environment and vice versa. The use of a conductive enclosure, e.g., a metal enclosure, means that the CSSP cartridge includes its own electromagnetic shielding. To this end the CSSP enclosure 121 is provided with EMI fingers 129 to ensure good contact with the shelf chassis and the adjacent components.

As shown in Figure 4, the CSSP cartridge 71 incorporates two D-shaped handles 132 to facilitate insertion and removal of the CSSP cartridge 71 with respect to an aperture 72 in the rear face of the shelf enclosure. A latch member 131 can be pivotably mounted on a plate that can be secured (e.g., using screws) to the rear face of the shelf enclosure. The latch member 131 is configured to engage one of the handles 132 and to secure the CSSP cartridge 71 in place. In other examples, the CSSP cartridge 71 could be provided with an injector/ejector handle in a manner similar to the information processing cartridge. As shown in Figure 6, the front face 123 of the CSSP cartridge 71 has perforations 133 to allow for airflow into the CSSP cartridge 71. As shown in Figure 4, the rear face 122 of the CSSP cartridge 71 has perforations 135 to allow for air to be exhausted from the rear of the CSSP cartridge 71.

20

At least one fan can be located, for example behind the perforated portion 135 of the rear face, in a CSSP cartridge 71 to channel cooling air through the CSSP cartridge 71 from the front to the rear. In this particular example shown, two fans are provided, one behind each set of perforations 135. LED indicators 137, as shown in Figure 4, can be provided on the rear face 122 of the CSSP enclosure 121 to indicate whether power is on, whether service intervention is required and whether the switch can be removed. Additional link status indicators can be provided integral to 2x4 stacked RJ-45 connectors 139, also shown in Figure 4. As shown in Figure 4, electrical connections 141 can be provided at the front face of the CSSP (i.e. on the

face that in use is inside the shelf enclosure 47). Suitable connections for use in the present example include a connector for power connections, a connector for serial management data connections and a connector for information connections. In the present example, information connections are implemented using an Ethernet
5 information communication protocol, e.g. at 1 Gigabit (Gb). However other protocols could equally be used, for example the Infiniband information communication protocol. The connector arrangement can include a guide pin arrangement to prevent module misalignment during insertion of the CSSP module into the receiving location. For this purpose, guide pin holes 142 can be provided on the front face 123
10 into which guide pins may pass to aid module alignment.

In the present example, up to two CSSPs 71 can be mounted at any one time at the rear of the shelf unit in corresponding apertures 72 in the rear face of the shelf enclosure 47. The number of CSSPs 71 provided in any particular implementation
15 depends upon system configuration and the need, or otherwise, for redundancy.

It will be appreciated that one possible construction of the CSSP cartridge 71 has been described and that as for the information processing cartridge 43, other examples could employ other materials and/or constructions.
20

Figure 7 provides a perspective view, partly from the front, of a power supply unit (PSU) cartridge 81. Here it is to be noted that the term "front" is applied in the context of the position, when installed, of the PSU cartridge 81, with respect to the shelf 41 (i.e. in this case the "front" of the PSU cartridge 81 is the innermost part of
25 the PSU cartridge 81 when it is inserted in the shelf 41).

With reference to Figures 4 and 7, it will be noted that PSU cartridge 81 is three-dimensional, having a height (h), width (w) and depth (d). In this particular example, the order to provide for a dense packing of the FRUs in the rear of the shelf

41, the PSU cartridge 81 has two dimensions (hereinafter described as the width, w, and the depth, d) that are generally similar. In a particular example, the enclosure of a PSU cartridge 81 has a height h, width w and depth d of 83mm, 202mm and 276mm, respectively, although of course, other examples may have different dimensions.

5

An enclosure 145 of present example of a PSU cartridge 81 is of generally oblong shape, but has the "top" "front" edge cut away to form an additional "top" "front" sloping face. The enclosure 145 therefore has five, generally rectangular, faces and two faces of generally rectangular shape with one corner cut away. For
10 convenience only, the face that is visible from the rear of the racking when the PSU cartridge 81 is mounted in the shelf 41 is known as the rear face 146. The opposite face is known as the front face 147. In the present example these two faces and the two side faces 150, 151 are of elongate, generally rectangular shape with one corner cut away, given that the width and depth of the PSU cartridge are similar, whereas the
15 top and bottom faces 148, 149, although still rectangular, are not, in this example, notably elongate. A top front face 148a is present at the top front of the enclosure. Thus the front of the enclosure is sloped at the top edge. As for the information processing cartridges 43, however, it will be appreciated that other examples could have other configurations.

20

In this example, the PSU cartridge enclosure 145 is fabricated from steel sheet to form a housing portion that includes the bottom face 149, the side faces 150 and 151 and the front and rear faces 146 and 147. Cover portions that are secured to the housing portion form the top face 148 and top front face 148a. The cover portions are
25 secured to the chassis portion by suitable fixings, for example one or more screws 152. It will be appreciated however, that in another example, other faces, or portions, of the enclosure could form the chassis and the cover portions. The provision of the enclosure 145 means that the PSU cartridge 81 can safely be handled by an operator who is not a skilled technician. Also, through the use of the enclosure 145, the PSU

cartridge 81 is a robust unit that protects its inner workings from the outside environment and vice versa. The use of a conductive enclosure, e.g., a metal enclosure, means that the PSU cartridge includes its own electromagnetic shielding. To this end the PSU enclosure 145 is provided with EMI fingers 153 to ensure good
5 contact with the shelf chassis and the adjacent components.

As shown in Figure 4, the PSU cartridge 81 incorporates two D-shaped handles 156 to facilitate insertion and removal of the PSU cartridge 81 with respect to an aperture 82 in the rear face of the shelf enclosure. A latch member 155 can be
10 pivotably mounted on a plate that can be secured (e.g., using screws) to the rear face of the shelf enclosure. The latch member 155 is configured to engage one of the handles 156 and to secure the PSU 81 in place. In other examples, the PSU 81 could be provided with an injector/ejector handle in a manner similar to the information processing cartridge. As shown in Figure 7, the front face 147 of the PSU cartridge 81
15 has perforations 157 to allow for airflow into the PSU cartridge 81. As shown in Figure 4, the rear face 146 of the PSU cartridge 81 also has perforations 159 to allow for air to be exhausted from the rear of the PSU cartridge 81.

A pair of fans can be located behind the perforated portions 159 of the rear
20 face of a PSU cartridge 81 to channel cooling air through the PSU cartridge from the front to the rear. LED indicators 161 can be provided on the rear face 146 of the PSU enclosure 81 to indicate whether input power is good, whether output power is good, whether service intervention is required and whether the PSU can be removed. Electrical connectors 163 can be provided at the front face of the PSU (i.e. on the face
25 that in use is inside the shelf enclosure 47) for connection to the shelf. The PSU 81 of the present example may suitably employ an SSI-MPS (Server Systems Interface - Midrange Power Supply) compliant right angle connector at the front face 147 of the PSU 81 to connect to the shelf 41. The power inlet 83 for each PSU 81 can incorporate a cable/connector retention mechanism (not shown) on the rear face 146

of the PSU to prevent accidental or malicious removal of the power input cord from the PSU 81.

In the present example, the shelf unit enclosure provides slots at the rear of the shelf unit for two hot-swappable, AC input PSUs 81. To provide redundant power supply, both PSU slots are populated. More generally, $N+M$ power supplies can be provided, where N is the minimum number of power supply units needed to support the components mounted in the shelf, and M is a selectable number of 0, 1 or more power supplies to provide for redundancy. In the present example, $N = M = 1$.

10

It will be appreciated that one possible construction of the CSSP cartridge 71 has been described and that as for the information processing cartridge 43, other examples could employ other materials and/or constructions.

15 If the full complement of information processing cartridges or switches are not fitted to the enclosure, then blanking panels/modules (e.g., the blanking panels 44 shown in Figure 3) are fitted to all unpopulated positions to maintain the correct airflow and thermal characteristics of the shelf, a safe internal operating temperature for the shelf life expectancy of all shelf components, electromagnetic compliance
20 (EMC) containment and electrostatic discharge (ESD) containment.

It will be noted that each of the FRUs, such as the information processing cartridges 43, is advantageously contained in its own robust enclosure to facilitate EMC containment, ESD containment, handling, storage and transportation. Each
25 FRU can be configured as a 'sealed' unit in the sense that it can be configured not to have field or customer serviceable parts internally. The FRUs can be configured readily to plug into the shelf enclosure and to be hot swappable. The FRUs can be keyed to prevent incorrect positioning and insertion into the shelf enclosure and are arranged positively to be retained in the shelf by a latching/locking mechanism.

The examples of FRUs described above are not provided with removable media. In the present example, internal data storage is provided by 2.5" IDE 9.5mm or 12.7mm profile hard disk drive (HDD) devices mounted internally in each information processing cartridge 43 and in the CSSP cartridge 71. The drives are not considered as FRUs and are not hot-swappable disk drives in the present example, although they could be in other examples. In other examples, the information processing cartridges can be configured without internal hard disk drives.

10 The internal configuration of the shelf 41 and a midplane 171 contained therein is described in the following with reference to Figures 8A, 8B 8C and 8D, and Figures 9A, 9B and 9C.

Figure 8A is a schematic plan view showing the internal configuration of an example of a shelf 41 with the cover 61 removed. Figure 8B is a schematic perspective view from above the rear of the chassis portion 47 of the shelf enclosure with the field replaceable units removed. Figure 8C is a schematic perspective view from below the front of the chassis portion 47 of the shelf enclosure with the field replaceable units and the base 51 removed. Figure 8D is a schematic perspective view from the front and above a part of the base 51 of the shelf 41. Figures 9A, 9B and 9C are, respectively, front, top and rear views of the midplane 171. In this example, the midplane is, in use, mounted vertically within the shelf 41 extending across the width W of the shelf 41 at a position approximately half way between the front and the rear of the shelf 41.

25

The vertically mounted midplane 171 extends, in this example, across the shelf 41 and allows for the electrical interconnection of the FRUs. The various apertures in the front and rear faces 57 and 59 of the shelf 41, in combination with the midplane 171, can be provided with guides (e.g., rails 181) and keying e.g., offset connector

positioning for the insertion of the FRUs into the enclosure and midplane 171. The midplane 171 can be a double-sided, or multi-layer printed circuit board (PCB) assembly that can be mounted vertically in a rigid manner within the enclosure. It can carry connectors 175 on a front surface 172 for making electrical connection with
5 corresponding connectors 120 on the information processing cartridges 43. It can also carry connectors 177 and 179 on rear surface 173 for making electrical connection with corresponding connectors 141 and 163 on the CSSPs 71 and the PSUs 81, respectively. Conductive tracks (not shown) on and through the midplane 171 can be provided to interconnect the various connectors. In addition, the midplane can
10 provide connectors for receiving corresponding connectors connected to first and second indicator boards 183 and 184 that each carry a respective set of LED indicators 69. In the present example, the midplane 171 is not configured as a FRU and is not hot swappable. It is perforated to facilitate airflow through the shelf 41. The midplane 171 can include openings 185, which cooperate with openings in the
15 enclosures of the FRUs 43 and 81, to provide a path for cooling air to pass from the front to the rear of the shelf 41, the cooling air being driven by fans in one or more of the FRUs, for example in the PSUs 81, possibly also in the information processing cartridges 43.

20 A plenum chamber floor member 94 can extend horizontally from the front of the midplane 171 to the front face 57 of the shelf enclosure, or chassis 47. The member 94 provides a floor for a plenum chamber 66, which is supplied with air via the apertures 65 in the front bezel and, in the illustrated example, the slot shaped aperture 68 in the front face 57 of the shelf enclosure 47. Although, for reasons of
25 ease of illustration a slot shaped aperture 68 is shown, a plurality of apertures 68 aligned with the blade receiving locations may be provided. The aperture or apertures 68 can serve both as air vents for a flow of air to the plenum chamber 66, and also as latching locations for latching portions at the top of the injector/ejector levers 111 for

the blades shown in Figures 5B and 5C. The top and sides of the plenum chamber are provided by the top cover 61 and side faces 53 and 54 of the shelf enclosure 47.

A plurality of cartridge guides 97 can be provided at the underside of the
5 plenum chamber floor member 94. In the present example, these guides comprise sprung wire members, e.g., of a resilient metal such as spring steel, that are attached to the top surface of the plenum chamber floor member 94 and extend through a plurality of apertures therethrough to result in a row of guides 97 at the underside of the plenum chamber floor member 94. This arrangement is shown in Figures 8B and 8C.
10 In Figure 8B, the sprung wire members 98 are shown attached to the top surface of the plenum chamber floor member 94. In the present example, the sprung wire members 98 are arranged in pairs, such that two guides 97 are provided by each spring clip 98. In Figure 8C, the guides 97 formed by the protruding portions of the sprung wire members 98 are shown at the underside of the plenum chamber floor member 94.
15 Each guide 97 is advantageously positioned so as to interface with the groove 108 in the plastics material in the upper face 104 of a processing cartridge 43 as shown in Figures 5A-5C to aid correct alignment and to facilitate insertion of the processing cartridge during insertion of the cartridge into the shelf 41. The use of the spring clip as a guide 97 also serves to urge the processing cartridge downwards to provide a
20 secure mounting of the processing cartridge 43, to take account of manufacturing and operational tolerances and to assist in insertion of the processing cartridge where an operator does not align this absolutely correctly.

A further row of cartridge guides 99 can be provided at the upper surface of
25 the base 51 of the shelf 41. In the present example, as shown in Figure 8D, these guides 99 have a rail like form, which can be achieved by punching or stamping through the base 51 of the shelf 41. In this example each guide, or rail, 99 includes a pair of upstands separated by an aperture 100 through the base 51. The size of the aperture 100 can correspond to the width between the upstands. The separation of the

upstands is selected so that the overall width of the resulting rails is slightly less than the width of a groove formed in the lower face of an information processing cartridge 43. Thus, each guide 97 is advantageously arranged so as to interface with the groove 1113 in the plastics material in the lower face 104 of a processing cartridge 43 as shown in Figures 5A-5C to aid correct alignment and to facilitate insertion of the processing cartridge during insertion of the cartridge into the shelf 41.

In the present example, where the guides 97 and 99 are formed from metal, the provision of the grooves 108 and 113 in plastics material at the upper and lower faces, respectively, of each information processing cartridge 43 (see Figures 5A-5B) results in a combination of metal and plastics materials that can provide a low friction interaction, facilitating insertion of the information processing cartridges.

If, for example, the information processing cartridge enclosure is made of a metal, it may be undesirable to provide metal guides to avoid a metal to metal interaction. In such a case, for example, it may be desirable to form the guides from a plastics material having a low coefficient of friction, such as polytetrafluoroethene (PTFE) or polythene. Plastics rails could be attached to the underside of the plenum chamber floor member 94 and/or on the upper surface of the base 51 of the shelf 41. In such an example, grooves on the upper and lower faces of the information processing cartridges 43 could then be formed of metal or plastics and still result in a low friction arrangement.

A CSSP/PSU divider 96 can be provided to the rear of the midplane 171 and can extend horizontally to the rear face 59 of the shelf enclosure 47. The CSSPs 71, when inserted, are supported by the divider 96. To aid the correct insertion of the CSSPs 71, CSSP guide pins 178 are provided on the midplane 171 at positions adjacent connectors 177 on the midplane 171 for connection to the CSSPs 71.

Respective positions 88 and 89 can be formed in the front face 57 and the rear face 59 at which first and second indicator boards 183 and 184 supporting the indicator LEDs 69 can be located. These positions 88, 89 therefore include an aperture through the respective face of the shelf enclosure 47 such that indicator LEDs 5 69 mounted onto a circuit board attached to the inside of the shelf enclosure 47 may be viewed from outside the shelf enclosure.

There now follows are more detailed description of the midplane 171.

10 As mentioned above, the midplane 171 connects all the elements of a shelf together, including, in the present example, up to sixteen information processing cartridges 43, up to two CSSPs 71, two PSUs 81 and the two indicator boards 183 and 184. In the present example, due to its location within the shelf enclosure, the midplane 171 is not configured to be swappable. Accordingly, to maximize the 15 system reliability, the midplane is configured to provide as a high level of reliability as possible. To this end, the midplane is advantageously configured without active devices and to include the minimum number of decoupling capacitors consistent with good design practice (ideally zero).

20 The midplane supports a number of paths for various power and signal lines to interconnect the FRUs.

In the present example, each information processing cartridge 43 has a high speed information signal connection (e.g., a Gigabit (Gb) Ethernet 25 SERializer/DESerializer (SERDES) connection) to each of the CSSPs 71, each connection consisting of two pairs of differential signals. In a conventional manner therefore, the tracking of the paths for these signals is arranged to keep the pairs well balanced and on a single signal layer (i.e. without vias) to support such differential signals at high frequency.

In addition, in the present example, each information processing cartridge 43 has a serial console connection to the CSSP cartridge 71. Each connection consists of two TTL (Transistor-Transistor Logic) level signals that make a transmit and return
5 (TX and RX) pair.

Also, each PSU 81 has a management signal connection (e.g., a serial I2C (Inter-IC Bus) connection) to the CSSP cartridge 71 to control power and monitor environmental parameters. The I2C bus comprises of two signals SCL and SDL
10 (serial clock line and serial data line). In addition, an I2C address programming pin is provided for the PSUs 81.

Each information processing cartridge 43 and PSU 81 can signal to the CSSP cartridge 71 that it is inserted by pulling to ground (GND) a respective Inserted_L
15 signal (i.e., an active low signal). These signals are fed to the CSSP cartridge 71 via the midplane 171.

Each PSU 81 has five 12 Volt output rails. The routing from each PSU 81 is arranged so that a fault in any single FRU cannot completely interrupt the power to
20 any other.

As mentioned above, the midplane 171 is provided with appropriate connector arrangements for receiving the connectors on the FRUs.

25 In the present example, the information processing cartridge 43 connects to the midplane 171 through a 40 pin Single Connector Attachment (SCA-2) connector as defined by the Small Computer Systems Interface (SCSI) standard. Accordingly, the midplane carries corresponding connectors 175.

In the present example, each CSSP cartridge 71 connects to the midplane 171 through a two right-angle 20 pair connector (e.g., 2mm HM-Zd connectors available from Tyco Electronics). The corresponding connectors 177 on the midplane are straight male parts with a power connector. A guide pin arrangement is provided in addition to the connectors to prevent misaligned modules causing bent pins during insertion. The guide pin also provides a leading ground. The CSSP cartridge 71 also connects to the midplane 171 through a right-angled 125 way 5 row 2mm connector. The connector 177 on the midplane 171 includes a straight male part. A guide pin arrangement is provided in addition to the connectors to prevent misaligned modules causing bent pins during insertion.

In the present example, as mentioned above, each PSU 81 connects to the midplane 171 through an SSI-MPS specification connector. The contacts are configured 5P/24S/6P with sequenced signal (S) and power (P) pins. Where the connector on the PSU is a 1450230-1 R/A male header, solder tails connector, the mating connector 179 on the midplane can be a 1450540-2 vertical receptacle, press-fit connector.

In the present implementation, indicator boards 183 and 184 (see Figure 8A) are provided at the front and rear of the system and are configured as FRUs. In this example they hold three system-level indicator LEDs 69 and include a FRU identity (FRU-ID) programmable read-only memory (PROM) each. Three LEDs 69 are present on the indicator board. There can, for example, be a white locator LED that can be switched by the user for locating the system; a green power-on LED to indicate when the system is powered; and an amber service-required LED to indicate a fault or other condition requiring servicing. These LEDs can be driven by the CSSP 71.

In the present example, identification information (FRU ID) for the midplane 171 is held on an I2C electrically erasable programmable read only memory

(EEPROM) in the front indicator board 183. In addition to the I2C signals necessary to access the FRU ID EEPROM, the CSSPs 71 provide a current limited supply to the indicator boards 183 and 184 via the midplane. The indicator boards 183 and 184 are also provided with an I2C address programming pin. Depending on the
5 implementation, FRU ID information can be stored instead, or in addition, on the rear indicator board 184.

As the FRU-ID for the midplane 171 is held on one or both of the indicator boards 183 and 184, the midplane can be a totally passive unit. The FRU-ID PROMs
10 communicate with the CSSPs 71 via an I2C bus. Each device on the bus has a separate I2C address. The lower three I2C address bits of the EEPROMs used are available as pins on the device, to allow programming with resistors. The least significant bit of this address (A0) is passed to the midplane via the corresponding connector. This allows the midplane 171 to program the address of the FRU-ID
15 differently for the front and rear indicator boards 183 and 184, by pulling the address low for the front board and high for the rear indicator board 183. This ensures that both EEPROMS are available on the bus, at different addresses. The FRU-ID for the midplane can be stored on either front or rear EEPROM, but the present example the FRU-ID is stored in the EEPROM on the front indicator board 183. The EEPROM
20 can be 8kByte or larger.

As mentioned above, the midplane 171 includes openings 185 to provide a ventilation path for cooling air passing through the shelf 41. The cooling air passing through the shelf 41 via the midplane 171 can be driven by means of fans provided in
25 each of the information processing cartridges 43 and the power supply modules 81. The openings 185 shown in Figures 8B, 9A, 9B and 9C form schematic representations of openings in the midplane 171. In practice, the openings could have any form (i.e., a series of large openings, or a number of small perforations), arranged on the midplane to align with corresponding openings or ventilation apertures in the

various field replaceable units 43, 71 and 81. In this way, the path of the airflow from the front of the shelf to the back of the shelf can be configured to be as efficient as possible, depending on the detail configuration of the fan units and the ventilation openings or apertures in the information processing, switch, service processor and power supply unit modules 43, 71 and 81. Providing the fan units in the field replaceable units 43, 71 and 81, contributes to the aim of maintaining the chassis 49 and the midplane 171 of the shelf 41 free of active components, thereby minimising cost, and facilitating maintenance. Also, by providing the fan units in each of the field replaceable units, merely inserting and removing field replaceable units automatically adapts the flow of cooling air to the number and type of field replaceable units inserted in the shelf 41.

As described above, in the present example each of the FRUs is designed to be a non-user serviceable unit. Thus each FRU presents the user with a "sealed" unit which may be inserted into and removed from the shelf 41 as desired or required. If a FRU ceases to be operable, then the user has a choice only of returning the FRU to a supplier or service company for repair or of discarding the non-operable unit. As the FRUs are non-user serviceable, there is no requirement for a skilled technician to be employed in inserting or removing the FRUs into or from a shelf 41. Thus each FRU is designed such that a non-skilled person should have difficulty in causing damage to the FRU during handling. Moreover, the configuration and construction of the FRUs (e.g., the injector/ejector levers, the grooves in the enclosures of the information processing units, etc), of the shelf enclosure and the midplane (e.g., the guide rails to guide insertion of the FRUs, the locating pins, etc) contribute to facilitating easy insertion and removal of the FRUs.

Shown in Figure 10 is an example of the flow of cooling air through the shelf 41 and FRUs 43, 71, 81 mounted therein.

In this example, the cooling air passing through the shelf 41 is drawn generally in a front to rear direction through the shelf 41 by cooling fans mounted within the CSSPs 71 and the PSUs 81. Two separate flow paths for cooling air are provided in this example. The first, indicated as flow path $\alpha\alpha$ by dotted lines 77 provides cooling
5 air to the CSSPs 71. The second path, indicated as flow path $\gamma\gamma$ by dotted lines 78 provides cooling air to the information processing cartridges 43 and PSUs 81.

The flow of cooling air along path $\alpha\alpha$ enters the shelf 41 through the aperture(s) 65 in the front face 57 of the shelf enclosure 47 into the plenum chamber
10 66. Further apertures (not shown) could also be provided in the side faces 53 and 55 of the shelf enclosure 47 to allow air to flow into the plenum chamber 66. This air then flows through the plenum chamber 66, and passes over the top edge of the midplane 171 to reach the perforations 133 of the front face of the CSSPs 71. The cooling air then passes through the CSSPs 71, providing cooling to the components
15 thereof before passing out of the CSSPs 71 through the perforations 135 in the rear face of the CSSPs 71 thus being exhausted from the shelf 41. This flow of cooling air along flow path $\alpha\alpha$ is driven by fans 79 mounted within the CSSPs 71. In the present example, a pair of fans 79 is provided within each CSSP 71 and is mounted against the rear face thereof.

20

Air flowing along path $\alpha\alpha$ is impeded from flowing around the processing cartridges 43 by plenum chamber floor member 94 and is impeded from flowing to the PSUs 81 by CSSP/PSU divider 96. This flow path $\alpha\alpha$ therefore ensures that air flowing to the CSSPs 71 is not warmed by passage through the processing cartridges
25 43 and therefore provides maximum efficiency cooling to the CSSPs 71.

The flow of cooling air along path $\gamma\gamma$ enters the shelf 41 through the perforations 115 in the front face of the information processing cartridges 43. The air thus enters the information processing cartridges 43 and provides cooling to the

components thereof. Cooling fans (not shown) within the information processing cartridges 43 direct the cooling air to the processor (CPU) of the information processing cartridge and direct the flow of air in the cartridge thereby increasing cooling efficiency. The air then exits the information processing cartridges 43
5 through the perforations 117 in the rear face thereof. The air then passes through the apertures 185 through the midplane 171 to reach the PSUs 81. This cooling air then passes through the perforations 157 on the front and upper front faces of the PSUs 81 to enter the PSUs and provide cooling to components thereof. It will be appreciated from Figure 10 that the sloping rear of the upper face of the PSUs 81 increases the
10 area over which air can be drawn into the PSUs, thereby reducing the back pressure on the air flowing through the shelf unit and aiding the cooling efficiency. The flow of cooling air along path $\gamma\gamma$ is driven by fans 85 mounted within the PSUs 81. In the present example, a pair of fans 85 is provided within each PSU 81 and are mounted against the rear face thereof.

15

Air reaching the PSUs 81 via path $\gamma\gamma$ will already have passed through the processing cartridges 43. Such air will therefore be already warmed above the ambient temperature outside of the shelf 41 by its passage through the processing cartridges 43. However, as the cooling requirement of the PSUs 81 is typically less
20 than that for the CSSPs 71, this does not cause any difficulty in the operation of the PSUs 81, which are adequately cooled by this flow of pre-warmed air. The pre-warmed air passing through the apertures 185 through the midplane 171 is impeded from flowing into path $\alpha\alpha$ and entering the CSSPs 71 by the SCCP/PSU divider 96.

25 As will be appreciated by the skilled addressee, the arrangement shown in Figure 10 is illustrative only and other arrangements whereby, for example, a mixture of cool air from the plenum chamber 66 and warmed air from the processing cartridges 43 is provided to each rear-mounted FRU can readily be constructed.

With reference to Figure 11, there now follows a description of functional elements of an information processing cartridge 43 as contained within the information processing cartridge enclosure 101.

5 The information processing cartridge 43 includes a microprocessor 192 (a non-limiting example of a microprocessor that can be utilised in the present example is an UltraSPARC™ processor). The microprocessor is mounted on an information processing cartridge motherboard 191.

10 A configurable clock generator 193, configured as a programmable clock synthesizer employing a crystal, can be used to produce CPU clock signals, CLKA and CLKB. The clock frequency can be determined by jumper settings (not shown). A vectored interrupt controller (I-Chip) 194 and a configurable core voltage regulator module (VRM) 195 are provided.

15 In the present example, memory means for use by the processor 192 when executing instructions can be provided in the form of buffered dynamic random access memory (DRAM), for example configured as dual in line memory modules (DIMMs) 196 with a 72-bit data path with error correction codes (ECC), seated in two sockets
20 on a riser card from the information processing cartridge motherboard 191. The memory capacity can be chosen to suit the processor addressable memory space. For example, in the present example, up to 4 Gigabytes (4GB) of addressable memory can be provided. Serial Presence Detect (SPD) auto-configuration is provided via a Service Management Bus (SMBus) over an I2C bus 197.

25 In the present example, a PCI bus architecture can be employed with a so-called SouthBridge bus bridge 199 with SuperIO and two Gb Ethernet Media Access Control (MAC) devices. As described above, however, other bus protocols (e.g., Infiniband) can be used. A 32bit PCI bus 198 can be provided from the

microprocessor 192. The SouthBridge 199 is a standard form of bus bridge, in the present example packaged in a 352 pin PBGA (Plastic Ball Grid Array) package, that provides the following functions: an SM Bus interface over the I2C bus 197 for access to the SPD (Serial Presence Detect) feature of the DIMMs that allows
5 initialization of the memory controller; an Xbus interface for access via an Xbus 200 (which is a packet switched multiprocessor bus) to a PROM 201, a real time clock (RTC) 202 and an information processing cartridge service controller (hereinafter termed a Blade Service Controller (BSC)) 203; an IDE (Integrated Drive Electronics) interface that provides an ATA-100 (AT Attachment) IDE connection 204 to an IDE
10 disk drive 205; and a serial console interface on a service bus 206 to the BSC 203 that is used for operating system functions including a console function with this embodiment.

For IO to the midplane 171, two AC-coupled Ethernet interfaces 207 and 208
15 are provided in the present example, which are packaged in a 316 pin PBGA. These Ethernet interfaces can provide a PCI attached Ethernet MAC capable of operation up to Gigabit Ethernet performance. The physical layer can be implemented using SERIALIZER/DESERIALIZERS (SERDESS) 209 and 210. An example of a SERDES device is the TLK2201 transceiver manufactured by Texas Instruments, Inc. The SERDES
20 devices use differential PECL TX+/- and RX+/- (Positive Emitter Coupled Logic Transmit and Receive) pairs to communicate to the switch portions of the CSSPs 71 over the midplane 171. The RX+/- pairs can be AC coupled at the information processing cartridge 43, the TX+/- pairs can be AC coupled at each CSSP 71. This facilitates hot-swap of the information processing cartridges 43 and the CSSPs 71.

25

Asynchronous serial connections 211 and 212 for communication between the BSC 203 and the Service Processor parts of the CSSPs 71 can be provided.

Internal data storage can be provided in the present example by a hard disk 205 with a capacity of 30GB or more rated for 24/7 continuous operation. The hard disk 205 is accessed using the primary IDE interface of the SouthBridge 199. The hard disk 205 can hold an operating system, for example a Solaris operating system, and
5 other software and data for performing information processing using the main, or host, processor (CPU) within the information processing cartridge 43.

In the present implementation, the BSC 203 can be implemented as a microcontroller (e.g., a Hitachi H8 microcontroller). The BSC 203 can provide
10 various functions, including for example: dual access (for the information processing cartridges and the CSSPs 71) to PROM 201 and EEPROM 213 for boot information and a FRU-ID for the information processing cartridge; channelling communication between an information processing cartridge 43 and the service processor part of the CSSPs 71; control of power on reset (POR), system reset and externally initiated reset
15 (XIR) to the microprocessor 192; control of the power, service-required and ready-to-remove LEDs 69; upgrading of field-upgradable firmware, via the serial interface; a watchdog function for the operating system; monitoring the speed of a CPU fan 214; and communications with an EEPROM 215 and the operating system via the Xbus 200.

20

In the present example, the BSC 203 can be powered by a 5V service bus (SB) rail as soon as a CSSP 71 and a PSU 81 are fully inserted into the midplane 171, it then turns on other DC/DC converters to provide power to the remainder of the information processing cartridge 43. A BSC reset signal can be derived from a simple
25 conventional power on reset (POR) generator that monitors a 5V supply rail.

In the present example a 1MByte Flash PROM 201 can be provided for storing boot variables for OpenBoot™ PROM (OBP) and Power-On-Self-Test (POST). Further OBP variables can be stored in a second 16kByte (16kB) I2C PROM 215,

accessible via the SouthBridge SM Bus port over the IC Bus 197. The PROM 215 can contain 8kByte for OBP variables and 8kByte of unused space. A 16kByte I2C EEPROM 213 that is accessible via the BSC 203 can contain BSC variables and FRU-ID variables. The EEPROM is nominally divided into 8kByte for FRU-ID and
5 8kByte for the BSC variables. Write protection for the FRU-ID is implemented by BSC firmware. Such write protection may be carried out by, for example, acknowledging instructions to write to the protected area, but not to carry out those write instructions.

- 10 An environmental monitor sensor 215 can be provided to monitor the CPU and ambient temperatures. This sensor can be accessible via the onboard I2C bus from the BSC 203.

The information processing cartridge 43 can be powered from two, diode
15 commoned, 9V power supply rails 216 and 217. DC/DC converters 218 can be used to provide the voltage levels required by the information processing cartridge 43. The DC/DC converters 218 are supplied by dual 9V inputs 216, 217, individually fused 219, 220 and then diode commoned 221, 222. A 5V DC/DC converter can be turned on as soon as the FRU is fully inserted, with the BSC 203 and required portions of the
20 SouthBridge 199 being powered (the 5VSB rail). A field effect transistor (FET) can be used to gate off the main 5V supply to the rest of the information processing cartridge 43. The DC/DC converter outputs and the main 5V FET can be arranged not to turn on until the BSC 203 turns them on via a signal from the SouthBridge 199. The SouthBridge 199 can be used so that if the BSC 203 is reset (by a watchdog
25 timeout or after a firmware download) the state of the DC/DC converters 218 is not affected. When the remaining outputs from the DC/DC converters 218 are within specification, a PWR_GOOD signal can be asserted low to the BSC 203.

A SouthBridge resume circuit can be operable to run from 3V3, and a simple Zener diode dropper circuit can be used to generate 3V3 from the 5VSB supply.

When the FRU is inserted the inrush current can be limited, for example to
5 <1A, and the rate of rise can be configured not to exceed a predetermined value (e.g., 20A/s) to provide a so-called soft start to facilitate hot-insertion. The intent is to prevent damage to the connectors and to avoid generating noise. A soft start controller 223, which controls a ramping-up of voltage levels, can be enabled when the predetermined signal (Inserted_L signal) is asserted low, this signal is on a short
10 pin in the connector and is connected to ground (GND – not shown) through the midplane 171.

In the present example, a processor impingement fan (processor fan) 214 is configured to run at full speed to cool the information processing cartridge 43 and the
15 fan. The speed of the processor fan and sink can be monitored by the BSC 203, using a tachometer sense pin on the microcontroller. In the event of the fan speed falling below a predetermined speed, or percentage of its nominal speed (e.g., 80%), the BSC 203 can be arranged to issue an alert. The nominal speed of the fan can be recorded as part of the BSC EEPROM contents.

20

The midplane connector 120 for the information processing cartridge 43 is used to establish the connection between the information processing cartridge 43 and the midplane. In the present example it supports up to 84 connections (pins) that will deliver SERDES outputs 224, 225, I2C signals 226, 227, and power 216, 217. Signal
25 connections may be made through a right-angled connector. Power connections may be made through the information processing cartridge right-angled connector. The connector can be configured to facilitate hotswapping of the information processing cartridge, for example by having a low insertion force and/or guide pins to increase the ease of serviceability and prevent module misalignment during insertion.

Interrupts to the processor 192 can be encoded using an encoded interrupt vector mechanism. An I-Chip Emulator (ICE) 228 functions as an interrupt concentrator, receiving all system interrupts and encoding them as an interrupt vector according to an interrupt vector code utilisable by the processor 192. In the present example, where an UltraSPARC™ processor is used, the interrupt vector encoding may be based on a 6-bit interrupt vector code.

With reference to Figure 12, there now follows a description of an example of a combined switch and service processor (CSSP) 71. In the present example, each CSSP 71 provides the functionality of a Switch 73 and of a Shelf Service Processor, or Shelf Service Processor (SSP) 74.

Figure 12 provides an overview of the functional components of the CSSP 71 including functional components of the Switch 73 and functional components of the SSP 74. In the present example, most of the components relating to the Switch 73 are mounted on a Switch PCB 231, and the components relating to the SSP 75 are provided on a SSP PCB 232. However, it should be noted that the components located in the lower portion of the switch PCB 321 (i.e., that portion below the SSP PCB 232 as illustrated in Figure 12 logically belong to the SSP 74, rather than to the switch 73. It will be appreciated that such component arrangements are not compulsory for successful operation and that any other component arrangement over any number of component boards can be easily achieved using conventional component arrangement techniques.

25

Firstly, with reference to Figure 12, there follows a description of functional elements of the Switch portions 73 of a CSSP 71 as contained within the CSSP enclosure 121.

The midplane connector 141 on the CSSP 71 establishes the connection between the CSSP 71 and the midplane 171. In the present example, it supports up to 84 connections (pins) that will deliver SERDES outputs 265-268, I2C signals 310, 320, 321 and 322, and power 278, 279. Signal connections may be made through two
5 20-pair right-angled connectors. Power connections may be made through a right-angled connector. The connector can be configured to facilitate hotswapping of the board, for example with a low insertion force. The connector also uses guide pins to increase the ease of serviceability and prevent module misalignment during insertion.

10 A switch microprocessor 240 is provided, in the present example the microprocessor used is a Power PC (MPC8245) packaged in a 352pin Tape Ball Grid Array (TBGA) package. This microprocessor 240 supports between 1MB and 2GB of address space in the present example. It further includes an Embedded Programmable Interrupt Controller (EPIC) that provides 5 hardware interrupts (IRQs) or 16 serial
15 interrupts. There are 4 programmable timers with cascade mode function. DRAM memory for the processor can be provided in the present example by a commodity DIMM 242. The processor 240 can be connected to a 32bit PCI bus 241, which operates at, for example, 33MHz/66MHz.

20 A clock input to the processor 240 can be provided by a clock generator (CLK) 243. The CLK 243 can include a configurable clock generator (not shown) implemented as a programmable clock synthesiser employing a crystal used to produce CPU clock signals. The clock frequency can be determined by jumper settings (not shown). A vectored interrupt controller (I-Chip) (not shown) and a
25 configurable core voltage regulator module (VRM) (not shown) can be provided that operate substantially as described above with reference to the like components of Figure 11.

In the present embodiment two switch ASICs (application specific integrated circuits) 244, 245 are provided (in the present example, BCM5632 Gigabit switch ASICs). Each ASIC can provide twelve GMII Interfaces (1 Gigabit Ethernet) (for uplinks and downlinks) and one 10Gb XGMII interface for chip-to-chip
5 communication (bridging) 246 between the ASICs 244 and 245. Sixteen GMII 1Gb 'downlinks', in the form of serialized Gb Ethernet data, are provided through four quad SERDES 248-251 to allow each information processing cartridge 43 to communicate with the switch 73. Eight GMII 1Gb 'uplinks' are provided for external communication through two quad PHYs 253 and 254 (in the present example
10 BCM5404 ASICs) and RJ45 connectors on the rear panel 122. The ASICs 244 and 245 are configured via a PCI interface (32bit/33MHz) to the PCI bus 241.

A Flash PROM 256 can store a real time operating system, and management and configuration data for the microprocessor. The Flash PROM 256 in the present
15 example can be operable to hold 8MB - 16MB of data, depending on the software required. The flash PROM 256 can be operated via an on-chip XBus 258.

Also connected to communicate with the processor 240 via the XBus 258, a Real Time Clock (RTC) 259 can be provided for real-time functions with a back-up
20 battery.

Also connected to the XBus 258 can be a UART (Universal Asynchronous Receiver Transmitter) 260 which in turn connects to a serial bus 261 for providing an asynchronous console connection from the switch 73 to the SSP 74 which can be
25 accessed by the SSP.

An integrated MAC/PHY (Media Access Control/Physical) switch 271 can provides its own interface to the PCI bus 241. This MAC/PHY switch 271 can connects to a 10/100 Ethernet hub 272. The hub 272 can be operable to provide a

management interface to the SSP 74 and a connection from an external management network to the switch 73 and SSP 74 of a given CSSP 71. The connection from the integrated MAC/PHY device 271 to the SSP 74 can be coupled capacitively. A loopback mode can be provided by the MAC/PHY device 271 for system diagnostics.

5 The hub 272 can connect to an RJ45 connector 273 on the rear panel 122 of the CSSP enclosure 121.

An 8kByte I2C EEPROM 262 can be used to store the FRU-ID and is accessible by the SSP portion 74 of each CSSP 71 via a serial bus 263 and the

10 midplane 171. The upper 2kByte of the EEPROM 262 can be configured to be write protected.

An I2C Redundant Control Register (RCR) 275 can be used to provide an alternate, redundant path for powering-down the CSSP 71 and Shelf Level Indicators

15 69 mounted on the front 57 and rear 59 panels of the shelf 41. The I2C RCR 275 can be accessible by both the SSP 74 of the CSSP 71 containing the RCR and the SSP 74 of a further CSSP 71 connected via the midplane 171 via an I2C bus 276. In the present example, a device suitable for use as the RCR 275 is a Phillips PCF8574 IC.

20 With continued reference to Figure 12, there now follows a description of functional elements of the Shelf Service Processor (SSP) portion 74 of a CSSP 71 as contained within the CSSP enclosure 121 and provided on an SSP PCB 232.

In the present example, communication between the Switch PCB 231 and the

25 SSP PCB 232 is facilitated by an interboard connector pair 298 and 299. It supports connections (pins) for I2C signals, 10/100 MAC/PHY output, and power. As described above, the switch PCB 231 carries the components associated with the switch, and it also carries the power, FRU-ID and environmental monitoring components along with the connectors for connections to the midplane 171 and

external connectors. Thus, in the present example, all SSP components requiring a connection to the midplane 171 or an external connection have signal paths routed through the connector pair 298, 299 and via the switch PCB 231 to the relevant midplane or external connectors.

5

In the present example, the SSP 74 includes a microprocessor 301 (e.g., a Power PC (MPC8245) processor) mounted on the SSP printed circuit board (PCB) 232. The processor 301 can be connected to a PCI bus 302, the present instance a 32 bit bus that operates, for example, at 33MHz/66MHz.

10

A clock input to the processor 301 can be provided by a clock generator (CLK) 303. The CLK 303 can comprise a configurable clock generator (not shown) implemented as a programmable clock synthesiser employing a crystal used to produce CPU clock signals. The clock frequency can be determined by jumper
15 settings (not shown). A vectored interrupt controller (I-Chip) (not shown) and a configurable core voltage regulator module (VRM) (not shown) can be provided that operate substantially as described above with reference to the like components of Figure 11.

20 The processor 301 can be provided with a DRAM memory 305. The memory capacity can be chosen to suit the processor addressable memory space. In the present example, 8 MB of DRAM memory is provided.

An integrated MAC/PHY switch 306 can provide its own interface to the PCI
25 bus 302. The MAC/PHY switch 271 can be connected to 10/100 Ethernet hub 272 via the interboard connectors 298, 299. A loopback mode can be provided by the MAC/PHY switch 306 for system diagnostics.

Octal UARTs 308 and 309 can be connected between the PCI bus 302 and the interboard connector pair 298, 299. The signal path can be continued from the interboard connector pair 298, 299 to serial connections 310 on the midplane connector 141 on switch PCB 231. The Octal UARTS 308, 309 can facilitate serial
5 communications between the SSP 74 and each of the processing cartridges 43.

Also connected to the PCI Bus 302 can be a dual UART (DUART) 312 that in turn can connect via the interboard connectors 298, 299 to serial bus 261 for providing an asynchronous console connection from the SSP 74 to the switch 73. The DUART
10 312 can also have an I2C connection to an external connector on the rear face 122 of the CSSP enclosure 121. The external connector can provide a common operating system/boot console and command port 311.

Connected to the processor 301 via an XBus 314 can be a Flash PROM 315.
15 The Flash PROM 315 can store a real time operating system, and management and configuration data for the microprocessor 301. The Flash PROM 315 can be operable in the present example to hold up to 2MB of data, depending on the software required.

Also connected to the processor 301 via the XBus 214 can be a real time clock
20 (RTC) 316 for real-time functions with a backup battery. The RTC 316 can also provide 8kByte of non-volatile random access memory (NVRAM), in the present instance implemented as an EEPROM. This can be used to contain information such as the FRU-ID, a serial number and other FRU information.

25 To facilitate I2C communications between the SSP 74 and the other CSSP 71, the midplane 171 and the PSUs 81, a multiplexer 318 can be provided. The multiplexer 318 can have a single I2C connection to the processor 301 and connections, via the interboard connector pair 298, 299 and the midplane connector 141 to both PSUs 81, the midplane 171 and the other CSSP 71.

The processor 301 can also comprise an embedded DUART to provide a redundant serial link to the SSP 74 of the other CSSP 71. Although it would be possible to implement this link using an external DUART, the advantage of using an
5 embedded DUART is that the connection to the other CSSP is reliable and therefore likely to be functional. Where the embedded DUART link does not use the I2C Multiplexer for communications to the other CSSP, a common mode of failure for both the SSP – SSP I2C links can be avoided, it being assumed that the processor 301 is likely to be functional even if both embedded DUART channels are non-functional.

10

The CSSP 71 can be powered from two, diode commoned, 9V power supply rails 278 and 279. DC/DC converters 281 can be used to provide the voltage levels required by the CSSP 71. The DC/DC converters 281 can be supplied by dual 9V inputs 278, 279, individually fused 285, 286 and then diode commoned 287, 288. A
15 soft start controller 283 can be provided to facilitate hot-insertion. A 5V DC/DC converter (I2C power regulator) 282 can be turned on as soon as the CSSP 71 is fully inserted. A 3.3V DC/DC converter can be turned on when instructed, for example through SSP service software, by asserting low an appropriate signal (ON_L – not shown). The 3.3V converter can be arranged to turn on a converted for 2.5V, 1.2V,
20 and a processor core voltage rail (Vcore) when the voltages are within an appropriate range.

When the CSSP 71 is inserted the inrush current can be limited, for example to <1A, and the rate of rise can be configured not to exceed a predetermined value (e.g.,
25 20A/s) to provide a so-called soft start to facilitate hot-insertion. The intent is to prevent damage to the connectors and to avoid generating noise. A soft start controller 283, which controls a ramping-up of voltage levels, can be enabled when the predetermined signal (Inserted_L signal) is asserted low, this signal is on a short pin in the connector and is connected to ground (GND – not shown) through the

midplane 171 until one of the supplies is removed. These circuits can be configured to withstand an overvoltage at their inputs whilst the input they are feeding is not powered, without any leakage to the unpowered circuit. A sense circuit can detect if the voltage has dropped below a threshold, for example 2.0V, as a result of a blown
5 fuse, a power rail going down, etc. The DC/DC converters 281 can be protected against short circuit of their outputs so that no damage occurs.

The I2C regulator 282 can be powered as soon as the CSSP 71 is fully inserted into the midplane 171. This can be facilitated through short pins connected to the soft
10 start controller 283, which controls a ramping-up of voltage levels. The other DC/DC regulators can be turned on, for example by SSP software.

A pair of fans 290, 291 can provide cooling to the CSSP 71. The fans 290, 291 can be configured to run at full speed to prevent overtemperature conditions by
15 minimizing the temperature of the internal components and the fan. The speed of the fans 290, 291 can be monitored by the SSP 74 through an environmental monitor 295 on the switch board 231. The environmental monitor 295 can be alerted in the event of the fan speed falling below a predetermined value (e.g., 80% of its nominal speed). The fan can provide tachometer outputs to facilitate the measurement of fan speed.

20

LED indicators 137 can be provided, for example with a green power LED, an amber LED for indicating that service is required and a blue LED for indicating that the switch is ready to be removed. LED indicators integrated on 2x4 stacked RJ45 connectors on the rear face of the CSSP 71 can be arranged, for example, to show
25 green continually when the link is present and flash green when the link is active.

The environmental monitor ENV MON 295 can be provided to maintain operational integrity of the CSSP 71. The ENV MON 295 can include limit values in limit registers and can monitor, for example, temperature within the CSSP enclosure

121, the CSSP power rails, including the 12V, 3V3, Switch Processor Core Voltage, CSSP Processor Core Voltage and the two 9V power feed rails 278, 279 from the midplane 171. The outputs of the DC/DC converters 281 can be fed in to A/D inputs of the ENV MON 295 for Watchdog comparisons to be made to the voltage limits set
5 in the limit registers. As noted above, the ENV MON 295 can also monitor the operating speeds of the fans 290 and 291. The ENV MON 295 can communicate with the SSP 74 of both CSSPs via an I2C bus 296.

For IO to the midplane 171 shown in Figures 8A-10, the midplane connector
10 141 can include sixteen 1Gb Ethernet connections 265-268 from four quad SERDES 248-251 and the I2C bus lines 596.

The SSP 74 can access the I2C devices (FRU-ID EEPROM, 8-bit I/O expansion chip, and the system hardware monitor) through the midplane 171.
15

For external IO, rear panel Gb Ethernet connections can be provided from the two quad PHYs 253, 254 to 2x4 stacked RJ45 connectors 139 (to give 8 uplinks). Each port can be an independent 10/100/1000 BASE-T (auto negotiating) port. The PHY devices 253, 254 can operate in GMII mode to receive signals from the 8-
20 Gigabit interfaces on the ASICs 244, 245.

The Power Supply Units (PSUs) 81 can configured such that when two or more PSUs 81 are connected in parallel in the shelf 41, failure of any one of the paralleled units shall not affect system operation. Moreover, one of the PSUs can be
25 installed or removed from a "live" system with or without input power applied. The outputs can have overcurrent protection.

The PSU can have an I2C interface to provide power supply status via the midplane 171. The PSU can have an internal temperature sensor that reports via the

I2C interface. The PSU fan speed can also be monitored and errors are reported via the I2C interface. Overvoltage and overcurrent sensors can also report via the I2C interface.

5 There now follows a description of aspects of an example of a power supply 81 with particular reference to Figure 13.

 When a power supply (e.g. mains power, or UPS type protected power) is connected to the cable connector 83, transformer, regulator and rectifier circuitry 400
10 can operate to generate a DC output (in the present example, 12V DC) from the input (in the present example 230/240V 50Hz AC or 110V 60Hz AC).

 In order to ensure good cooling reliability within the shelf 41, each PSU 81 can have a pair of cooling fans 402, 403 located at the rear of the PSU enclosure as
15 described above with reference to Figure 7. As a failsafe measure against a failure of one PSU 81, the fans of each PSU 81 can be powered by both PSUs 81. Thus, in the event that one PSU 81 is non-operative to produce the DC supply for the shelf 41, as long as the other PSU 81 is operative, not only will all components of the shelf 41 be powered as normal, the fans of both PSUs 81 can continue to run.

20

 As shown in Figure 13, this dual powering of cooling fans 402, 403 can be effected by providing a power supply line 404 from the transformer, regulator and rectifier circuitry 400 to power both fans 402, 403. Also, first and second separate power lines 410, 412 from the other PSU 81 can provide duplicate power supply to
25 the first and second fans 402, 403 respectively. The fan 402 can thus be powered by a diode commoned supply from line 404 and a diode commoned supply from line 410. Diode protection can be provided by diodes 405 and 411 respectively. The speed of the fan 402 can be controlled by a speed controller 408. Similarly the fan 403 can be powered by a diode commoned supply from line 404 and a diode commoned supply

from line 412. Diode protection can be provided by diodes 406 and 414 respectively. The speed of the fan 403 can be controlled by a speed controller 409.

The two speed controllers 408, 409 can in turn be controlled by a data input
5 from each CSSP 71 received via an I2C bus connection (not shown in Figure 13). Power supply lines carrying DC power for the other FRUs of the shelf 41 are shown in Figure 13 as power line 416. All power connections to and from the PSU 81 can connect to the midplane 171 when the PSU is inserted in the shelf 41 via the midplane connector 163. In the present example the PSU 81 connects to the shelf through a
10 5P/24S/6P configuration SSI-MPS compliant right angle connector 163 at the front face 147 of the PSU 81. Connectors for the I2C interface can also be provided.

The input power line 410 and 412 for each fan 402 and 403 can be provided with a softstart module 4131 and 4132 respectively, to allow for hot insertion of the
15 PSU 81 into the shelf 41. The softstart modules 4131 and 4132 can be controlled, for example, by pulling a signal to ground (e.g., a "mated" input line 4151 and 4152).

Where the two input power lines 410 and 412 are separate lines having separate softstart provision, there is no common failure mode for the backup method
20 of powering the fans 402, 403. Thus even if the PSU 81 ceases to be operable to generate the DC supply, and a component (power line or softstart module for example) fails in the supply path from the second PSU 81 to the fans 402, 403, at least one of the fans 402, 403 can remain operational as the shelf 41 still receives the cooling effect of three PSU fans.

25

In the present example, the power supply has four rear panel LED indicators 137. A blue "Ready to Remove" LED can be driven by the I2C interface and indicate that the power supply may be removed from the system. An amber "Service Required" LED can be driven by the I2C interface and indicate that the power supply

is in a fault condition: any output out of range, over-temperature or shutdown. A green "DC Output-OK" indicator can be driven by internal power supply circuits and show that the main 12 volt supply is functioning. The LEDs can remain lighted when individual outputs are in the current limited mode of operation. A green "AC Input-
5 OK" indicator can be driven by internal power supply circuits and show that AC input power is within normal operating range.

With reference to Figure 14, there will now be described an example of data connectivity between the FRUs and midplane of the shelf 41. Power transmission
10 paths are not illustrated in Figure 14. However, it will be appreciated that to facilitate the maximum component redundancy of the shelf 41, each PSU 81 can independently provide power to each FRU.

In the present example each of the processing cartridges (blades) 43 connects
15 to the midplane 171 via a pair of information signal connections (e.g. Gb Ethernet links) 224, 225 and a pair of serial management signal connections 226, 227. Connections within the midplane 171 can ensure that each Ethernet link 224 is directed to a connection 265-268 from the midplane 171 to a first switch 73, and that each Ethernet link 225 is directed to a connection 265-268 from the midplane 171 to a
20 second switch 73. Thus one Ethernet link can be established between each processing cartridge 43 and the switch 73 of each CSSP 71. Further connections within the midplane 171 can ensure that each serial connection 226 is directed to a connection 310 from the midplane 171 to the first SSP 74 and that each serial connection 227 is directed to the second SSP 74. Thus one serial link can be established between each
25 processing cartridge 43 and the SSP 74 of each CSSP 71. As mentioned earlier, information signal connections other than Gb Ethernet connections (e.g. Infinband connections) could be employed in other examples.

A plurality of serial connections can connect each SSP 74 to the other. Serial lines 320, 321 can connect each SSP 74 to the midplane 171 and connections within the midplane can connect the two sets of lines together. To provide a control interface from the SSPs 74 to the PSUs 81, serial lines 322 can connect each SSP 74 to the midplane 171 and connections within the midplane 171 can connect to serial lines 324 from the midplane 171 to each PSU 81.

An example of data and control connectivity of the shelf 41 to and from computer systems external to the shelf 41 when the shelf 41 is arranged for use within a multiprocessor server system such as that described above with reference to Figure 1 will be described with reference to Figure 15.

As summarised above with reference to Figure 14, in the present example each processing cartridge, or blade, 43 is connected to the switch 73 of each CSSP 71 by an information signal connection (e.g. a 1Gb Ethernet link) formed by a combination of links 224, 225 from the processing cartridge 43 to the midplane 171, connections within the midplane 171 and links 265-268 from the midplane 171 to each switch 73.

Further, in this example a set of serial management signal connections comprising links 320, 321 and connections within the midplane 171 connect the SSP 74 of each CSSP 71 to the SSP 74 of the other CSSP 71.

To provide external data connectivity between the shelf 41 and an external core data network 330, in association with which all information processing performed by the processing cartridges 43 of the shelf 41 is undertaken, connections 331 can be formed between the core data network 330 and the eight 1Gb Ethernet ports 139 provided on the rear panel 122 of the CSSP enclosure 121.

In the present example, the connections by means of which control and configuration of the shelf 41 are performed are entirely separate to the connections to the core data network 330. Therefore, a first external switch 335 can connect to a management (I2C) port 273 of the first CSSP 71 and a second external switch 336 can
5 connect to a management (I2C) port 273 of the second CSSP 72. As described above with reference to Figure 12, the management port 273 can provide a management network interface to both the switch 73 and SSP 74 of each CSSP 71. The external switches 335, 336 can each be connected to each of a pair of System Management Server (SMSs) 338, 339. The SMS is not essential to the operation of the shelf 41,
10 but use thereof aids optimal operation of the shelf 41. In a typical multiprocessor server system a plurality of shelves 41 may be connected together via the core data network 330 under the control of a single management network utilising one set of SMSs 338, 339. A set of SMSs 338, 339 may comprise a single SMS (as well as a plurality thereof). However use of at least two SMSs enables redundancy of
15 components, therefore increasing overall system reliability.

A serial interface control 343 operable under telnet protocol control is also connected to the shelf 41 in the present example. This can provide a common operating system/boot console connection to the SSP 74 of both CSSPs 71 via the
20 RJ45 connector 311 on the rear panel 122 of each CSSP enclosure 121.

It will be appreciated from the above that a flexible and scalable modular computer architecture has been described. In the described example up to 16 information processing cartridges, or blades 43, can be configured as sealed FRUs on
25 a single shelf 41, the number of blades being chosen according to customer requirements. Each blade has its own processor and random access memory. If, for example, there is a maximum of 2Gbytes of memory per information processing cartridge, and one processor per blade, 16 processors (16P) with 5.33 processors per unit height (1U) and a total of 32GB of memory per shelf can be provided.

In the present example, the shelf 41 incorporates redundant combined switch and shelf service processor modules (CSSPs) 71 and redundant power supply units (PSUs) 81 separate from the blades 43. As the power supplies are carried by the shelf,
5 the information processing cartridges can be kept compact and inexpensive. Also, as a result, they can be powered by DC power only, via the midplane 171.

Also, as mentioned earlier, the FRUs (e.g., the information processing cartridges, or blades, 43, the CSSPs 71 and the PSUs 81) can all be configured as
10 sealed units that do not contain any internal FRUs themselves and do not contain user serviceable items. The enclosures of the FRUs can be arranged to enclose all of the functional components of the FRU with only electrical connectors being externally accessible and with indicator LEDs being externally visible as well.

15 These factors can all contribute to keeping the cost of the FRUs low, as well as that of the overall system. The modular approach with the use of sealed modular field replaceable units for providing system functionality and with non-field replaceable units designed with a minimum possible number of active components enhances reliability. Moreover, easy and rapid maintenance is facilitated in the event of a
20 failure of a FRU by simple replacement of that FRU, further reducing the cost of ownership.

Thus, it will be appreciated from the above description that the provision of a rack mountable shelf, that includes power supplies, a shelf service processor and
25 switches in modular units, for carrying a number of processing cartridges, wherein the number of processing cartridges can be chosen according to customer requirements, provides a flexible and scalable computer configuration. The balancing of the load between the processors of the processing cartridges can be effected by software using conventional principles.

A configuration as described provides an easily scalable processor architecture, whereby the processing power provided by a complete system based on the information processing cartridge/information processing cartridge carrier architecture can be scalable from moderate to very high capacity through the simple addition of further information processing cartridges.

As illustrated in Figure 16, an example of the external connections from a shelf 41 can be in the form of two active information signal connections (e.g., Ethernet connections) 350 and 351, two active power connections 353 and an active/standby pair of management connections 354. With regard to the management connections, each connection comprises a serial connection and a network (e.g. Ethernet or Infiniband) connection. It is possible to connect to either the active or the standby connection, as the incoming signal will be internally routed to whichever management controller (CSSP) is the current master. It will be appreciated, therefore, that the connections to a shelf can be kept to a minimum. It will further be appreciated from the configuration shown in Figure 16 that the system is scalable beyond a single shelf unit 41.

Figure 17 illustrates how a plurality of shelves can be configured within one (or more) racks to provide even higher processing power. Such a constellation of shelves to provide a large grouping of servers is sometimes termed a "web farm" or "server farm" 360. As shown in Figure 17, the web farm comprises a plurality of shelves 41 that each carry a plurality of blades 43. Also provided are a plurality of Network Attached Storage devices (NAS) 373 for providing storage for critical data, e.g. email data storage, for the web farm. The NASs 373 are not required if there is no critical data to be stored, e.g. if the web farm is operating solely to provide web caching services.

Management control of the web farm 360 can be provided through a pair of System Management Servers (SMSs) 362. Each SMS 362 can be connected to a management network via a link 366 and to a management console 365. The SMSs 362 can communicate with the individual shelves 41 via a pair of management switches 364. Each shelf 41 and NAS 373 can be connected to each management switch 364 via a connection 367. Thus dual redundant management connections can be provided to each shelf 41 and NAS 373.

Flow of data to and from the web farm 360 can be provided through a pair of data switches 369. Each data switch 369 can be connected to a consumer network via a link 370. It is to be understood that the consumer network can be a larger data network to which the web farm 360 is connected. This network can be an office or corporation intranet, a local area network (LAN), a wide area network (WAN), the Internet or any other network. Connections between the data switches and the shelves 41 can be facilitated by connections 371. It is to be noted that as each shelf has its own switching capability, there is no need for each shelf 41 to be directly connected to the data switches 369. Connections can also be provided to connect the NAS units 373 to the shelves 41. The topology used for interconnection of the data switches 369, shelves 41 and NASs 373 can be any topology providing at least one connection of any length between every possible pair of units. Complex topologies arranged to minimise the maximum connection length between any two given units in the web farm can be used.

The web farm 360 comprising a plurality of shelves 41 with or without a plurality of NASs 373 can suitably be used as any or all of the entry edge server group 9, web edge server group 15 and application servers 19 described above with reference to Figure 1.

As an alternative to providing critical data storage within a NAS 373, such storage can be provided within one or more NAS cartridges fitted into one or more of the shelves 41 in place of processing cartridges 43. Another alternative is to provide a server shelf with local storage such as a RAID array (Redundant Array of Inexpensive
5 Disks) in place of the NAS 373.

Thus there has now been described an example of a fully configurable computing system based on a plurality of self contained field replaceable units (FRUs) and scalable from a single processing cartridge with power supply and switching
10 capability to a multiply redundant multiprocessor server system with full system management capability extending over a number of co-operably connected server shelves. It will of course be readily apparent to the skilled reader that many of the specific features specified in the above description are in no way limiting and a variety of alternatives may be produced using only ordinary skill and common general
15 knowledge. Non-limiting examples of example modifications which may be made to the above described system are discussed hereafter.

There is no limit placed on the processing cartridges as to what software they should run. Each module within a shelf or farm may run under the same operating
20 system, or a plurality of different operating systems may be used. Examples of possible operating systems include Sun Microsystems' Solaris® OS or another UNIX™-Type OS such as Linux™, MINIX™, or Irix™, or UNIX™ or a Microsoft OS such as Windows NT™, Windows 2000™, Windows ME/98/95™, Windows XP™.

25

It is also not necessary that each processing cartridge within a shelf or farm be configured to run the same program software. For example, individual processing cartridges may be configured to execute, for example, fileserver software, mailserver

software, webhosting software, database software, firewall software, or verification software.

Although it has been described above with reference to Figures 4, 8, 12, 13
5 and 14, that functionality of a switch and of a shelf service processor should be provided within a single combined switch and service processor unit, this is not essential and separate switch and shelf service processor field replaceable units may be used.

10 Although it has been described above that a pair of PSUs and a pair of CSSPs may be provided so as to enable dual-redundancy, further PSUs and CSSPs may be provided so as to increase FRU redundancy further, thus providing statistically higher reliability.

15 In the power supply circuitry in each of the blades and CSSPs, two voltage sense circuits may be provided after the fuses and before the diodes, to prevent a latent fault caused by a failed fuse going undetected until one of the PSUs is removed or taken offline. Such circuits may configured to withstand an overvoltage at their inputs whilst the input they are feeding is not powered, without any leakage to the
20 unpowered circuit.

Although it has been described above with particular reference to Figure 11 that the processing module may be based on an UltraSPARC™ processor, this is not limiting and any other processor having sufficient processing capacity to undertake the
25 tasks required of a particular processing cartridge may be used. Alternative processors include, but are not limited to, Intel x86 series and compatible processors, AMD x86 compatible processors, Alpha processors and PowerPC processors. The particular example of an x86 compatible processor is described in more detail with reference to Figure 18. In Figure 18, the parts corresponding to those of the UltraSPARC™ based

system of Figure 11 have the same reference numerals and will not be described again here. In a system based on an x86 compatible processor, the processor 378 itself communicates with the other components, including the memory 196 and PCI bus 198 via a Northbridge 379. The Northbridge 379 also includes an interrupt controller, so
5 no separate interrupt concentrator is required. The other components of the processing cartridge could be substantially the same as for the UltraSPARC™ based system described above.

Although it has been described above that each information processing
10 cartridge comprises a single microprocessor, this is not a limiting case as each or any of the information processing cartridges may have more than one microprocessor arranged to share common storage resources to operate synchronously (in lockstep) or asynchronously. Also, it is not necessary that all information processing cartridges inserted into a shelf at a given time are identical, rather a variety of different blade
15 architectures may be used simultaneously.

The provision of the functions of both Switch and Shelf Service Processor within a single FRU in the present example provides a facility within a single shelf 41 for dual redundancy in both functions in fewer different FRUs. As will be
20 appreciated, there is no restriction that these two functions must be provided within a single FRU and division of the two functions into separate FRUs would present no difficulty to the skilled addressee.

The backup power provision for the cooling fans of the PSUs is described
25 above with reference to Figure 13. Although it is described that the backup power supply to each fan should be independent of the other, if the risk of common mode failure of backup power supply is judged to be low, or is of low importance, then the backup supply to each fan may be provided in common with the backup supply to all other fans.

Also, an information processing module forming a field replaceable server blade can include a processor and memory can be configured by means of software, firmware or hardware to provide a special purpose function. By way of examples
5 only, an information processing module can be configured to perform the function of one or more of a firewall, or a load balancer, encryption and/or decryption processing, an interface to a secure network, e.g. a virtual private network (VPN), a specialized switch with wide area network (WAN) connectability.

10 Also, a storage blade may be provided. The storage blade can be configured to be mountable in a server blade receiving location in a blade server carrier. The storage blade can comprise storage blade connectors configured for interconnecting with carrier connectors on the server blade carrier, whereby the storage blade is interchangeable with a server blade. A carrier, or shelf, for a server system can be
15 arranged with a plurality of blade receiving locations for receiving blades, wherein the blades can be storage blades or information processing blades. The server system can be self configuring on receipt of the blades according to the type of blade received in each said location. To achieve the blade service controller in each blade can be operable to communicate with a shelf service processor to perform said configuring.

20

One aspect of ensuring maximum possible computer availability relates to service and management of a computer system. In particular, service and management functions are in general required for monitoring system status and performing control actions in response to monitored parameters. There now follows
25 an example of how service functions may be organised to be implemented in a hierarchical manner through different components of a modular computer system.

As described above with reference to Figures 11 and 18, each information processing cartridge 43 has a blade service controller 203. The blade service

controller 203 is a local controller of service functions for each information processing cartridge 43. Each blade service controller 203 is arranged to be able to communicate with the SSP 74 (see Figures 12 and 15) of each CSSP 71 via the connections 226, 227 and 310. The SSP 74 provides shelf-level service functions for the population of each shelf 41, which includes the PSUs 81 and the switches 73 of the CSSPs 71 as well as the processing cartridges 43. Each SSP 74 is arranged to be operable to communicate with an external system management server (SMS) 338, 339 via I/O connector 273 as illustrated in Figure 15. The system management server provides system level service functions for one or more shelves 41. The system management server may also provide service functions to computer system modules other than shelves 41 within a computer system, for example to individual server units and/or to reliable storage units such as Network Addressable Storage (NAS) units.

In the present example, the blade service controller 203 can provide the following service functions. It typically performs environmental monitoring of the information processing cartridge 43 within which it is situated, this includes monitoring of CPU and ambient temperatures by means of the environmental monitor 215, monitoring of the power supply rails from the DC to DC converter 218 and monitoring of the speed of the processor cooling fan 214. The blade service controller 203 may also perform a watchdog function for monitoring the operating system running on the processor of the processing cartridge 43. The blade service controller 203 can also provide an out-of-band interface between the processing cartridge 43 and the CSSP 71. The out-of-band interface is a physically discrete interface which is independent of the operating system and applications running on the processor, this interface may be used for providing processor cartridge parameters to the shelf SSPs 74 at boot. In addition, the blade service controller 203 may also provide local functions for control of the LED service indicators 119 and control of power on reset (POR) and externally initiated reset (XIR) to the processor. In the present example, there is no requirement for the blade service controller 203 to have any level of

intelligence, it being sufficient that it is capable of performing simple watchdog and threshold type monitoring operations. However, the blade service controller 203 may be implemented having intelligence, which may be of particular benefit in the event that it is tasked with more complex functions than those set out above.

5

The SSP 73 (of which two are typically provided within a shelf 41 to provide redundancy) provides the following functions in the present example. It typically performs environmental monitoring for the shelf 41, monitoring the PSUs 81, the CSSPs 71 and the PSU inputs. The SSP 73 can also aggregate data from the blade
10 service controller 203 of each blade 43 present in the shelf. The SSP 73 may also be configured to make decisions regarding aspects such as power control, control of the shelf service indicator LEDs 69 and to provide instructions to the blade service controllers 203. In the present example, the SSP 74 is not responsible for deployment of software to individual processing cartridges, however this function may be
15 provided by the SSP 74 in alternative arrangements.

In order to provide a facility for manual monitoring and control of the SSP 74, it may be configured to have a user interface via which a user can be provided access to the functions and data of the SSP 74. The user interface access may be provided
20 through one of the external serial interface connection 311 and the external network connection 273. The SSP 74 may be operable to operate as an isolated unit with no higher level service control provided. In general the SSP 74 requires little or no user intervention to manage the operation of a single shelf 41.

25 Fault management may be controlled by the SSP 74 such that a failed processing cartridge 43 (the failure of which may be detected, for example, from data provided by a blade service controller 203 indicating a failure of the processing cartridge's operating system or application or from errors detected in the switch 73) may be remotely restarted by the SSP 73, in the present example either by initiating an

externally initiated reset (XIR) to the processor of the processing cartridge, or by causing an interruption in the power supply to the processing cartridge sufficient to cause a power on reset (POR) to the processor of the processing cartridge. In the present example, a failed processing cartridge will be the subject of a predetermined
5 number of restart attempts before that processing cartridge is considered permanently failed, at which time the processing cartridge is removed from service in that it is ignored by the SSP 74 and the switch 73 is instructed to pass traffic to it and to ignore any traffic from it and the failed nature is reported to any higher level service control to which the SSP 74 is connected.

10

In order to perform the above described functions, the SSP 74 of the present example is configured to have a degree of intelligence such that decision making steps and provision of a user interface are supported.

15 The system management server (SMS), of which two (338 and 339) are typically provided as a clustered pair for redundancy (as illustrated in Figure 15), is in the present example configured to provide the following functions. The primary use of the server management system 338, 339 of the present example is to provide aggregate management and monitoring for a number of shelves 41 (as illustrated in
20 Figure 17). By means of such aggregated monitoring and management, a single (clustered pair of) system management server(s) can oversee the running of a large number of computer system shelves 41.

The system management server (SMS) 338, 339 may also be operable to
25 perform hardware fault failover in response to hardware failures detected by the system management server or reported to it by an SSP 74 of a particular shelf 41. In addition, the system management server 338, 339 of the present example may be operable to perform software deployment to individual processing cartridges. In

alternative arrangements, the system management server 338, 339 may be operable to allow platform specific abstraction.

By platform abstraction is to be understood that the SMS 338, 339 may be
5 operable to allow platform specific abstraction such that the user may configure computing services without knowledge of the underlying hardware providing these services. For example a user may require a web server which can cope with, say, 1000 connections per an hour. The SMS 338, 339 may be pre-configured with enough platform specific data to determine how many blade servers are required to
10 fulfil this requirement, based on their processor, memory size, etc and provision the web service on that number of blade servers without the user needing to know anything about the server hardware details himself.

The operation of the system management server 338, 339 may be controlled
15 via a management network 366, such that control may be automated (via a higher level of system management) or manual (via a terminal to the management network 366) or a mixture of both. The system management server 388, 399 may also be directly connected to a management terminal 365 for direct operator control of the particular system management server. In order that the above functions be
20 implemented in the system management server, the system management server may be configured to have a level of intelligence appropriate to the making of decisions and controlling of failover and deployment.

Thus there has been described an arrangement wherein management and
25 service functions are distributed in a hierarchical manner across various levels of a modular computer system. Through the implementation of a hierarchical arrangement of service functions, any given service function may be implemented at a most appropriate level. For example, a top level service controller need not be concerned with environmental monitoring within an individual processing cartridge, whereas the

same top level service controller may have responsibility for overall management for a plurality of shelves 41. As the skilled addressee will appreciate, a number of alterations and modifications may be made to the above described arrangement without departing from the scope of the present invention. Also, a number of the
5 above described components and functions may be replaced with or supplemented by equivalent components and functions without departing from the scope of the invention.

Although the embodiments above have been described in considerable detail,
10 numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such variations and modifications.

CLAIMS

1. A server system comprising a server blade carrier carrying a plurality of removably mountable server blades, each blade including a blade service controller operable to perform monitoring and management functions in the blade, the carrier comprising at least one carrier service processor operably connected to the blade service controllers of blades mounted in the carrier to provide higher level management functions and configured to communicate management functions to and from the blades.
2. The server system of claim 1, wherein the blade carrier comprises at least two carrier service processors.
3. The server system of claim 2, wherein said at least two carrier service processors are operable in a redundant manner.
4. The server system of claim 3, wherein a first service processor is operable as a master service processor and a second service processor is operable as a backup service processor.
5. The server system of any preceding claim, wherein the carrier comprises a power supply unit operably connected to supply power to the carrier service processor and to the blade servers mounted in the carrier and wherein the carrier service processor is configured to provide management functions to the power supply unit.
6. The server system of any preceding claim, wherein the carrier comprises a switch operably connected to perform distribution of information signals within the carrier and wherein the carrier service processor is configured to provide management functions to the switch.

7. The server system of claim 6, wherein the carrier service processor and the switch are arranged into a unitary combined switch and service processor module.
- 5
8. The server system of any preceding claim, wherein the blade service controller is operable to monitor an environmental variable of the blade comprising the blade service controller.
- 10 9. The server system of claim 8, wherein the environment variable includes at least one of a processor temperature, an ambient temperature, a voltage level and a fan rotation speed.
10. The server system of any preceding claim, wherein the blade service controller is operable to monitor the operation of an operating system of the blade.
- 15
11. The server system of any preceding claim, wherein the blade service controller is operable to provide an out of band interface between the blade and the carrier service processor.
- 20
12. The server system of any preceding claim, wherein the blade service controller is operable to control a processor reset function.
13. The server system of any preceding claim, wherein the carrier service processor is operable to monitor an environmental variable of the blade carrier.
- 25
14. The server system of any preceding claim, wherein the carrier service processor is operable to cause a restart of a blade in response to a blade fault.

15. The server system of any preceding claim, wherein the carrier service processor is operable to cause a blade to be removed from service in response to a blade fault.
- 5 16. The server system of any preceding claim, wherein the carrier service processor is operable to report the removal from service of the blade to a higher level management provider.
- 10 17. The server system of any preceding claim, comprising at least one blade carrier and a system management server, the system management server being operably connected to the carrier service processors of the blade servers mounted in the carrier to provide high level management functions.
- 15 18. The server system of claim 17, comprising at least two blade carriers.
19. The server system of claim 17 or claim 18, comprising at least two system management servers.
- 20 20. The server system of claim 19, wherein said at least two system management servers are operable in a redundant manner
21. The server system of claim 19, wherein said at least two system management servers are operable as a clustered group.
- 25 22. The server system of any of claims 17 to 21, comprising a management network, the management network being operably connected to the system management server to provide high level management functions.

23. The server system of any of claims 17 to 22, comprising a management terminal, the management terminal being operably connected to the system management server to provide high level management functions.
- 5 24. The server system of any of claims 17 to 23, wherein the system management server is operable to perform deployment of software to a blade.
25. The server system of any of claims 17 to 24, wherein the system management server is operable to perform hardware fault failover of blades.
- 10 26. The server system of any preceding claim, wherein the blades are configured to be hot-swappable.
27. The server system of any preceding claim, wherein the blades are configured
15 as field replaceable units.
28. The server system of any preceding claim, wherein the carrier service processor is configured as a field replaceable unit.
- 20 29. The server system of any preceding claim, wherein the blade carrier is a rack-mountable shelf.
30. A method of providing management functions within a server system having a server blade carrier carrying a plurality of removably mountable server blades,
25 the method comprising:
performing monitoring and management functions at a blade service controller within each blade;
performing higher level management functions at at least one carrier service processor within the carrier;

providing a communications channel connecting the carrier service processor to each blade service controller; and
communicating management functions between the carrier service processor and the blade service controllers.

5

31. The method of claim 30, further comprising:
performing higher level management functions at a system management server remote the carrier;
providing a communications channel connecting the system management
server to the carrier service processor; and
communicating management functions between the system management server and the carrier service processor.

10

1/25

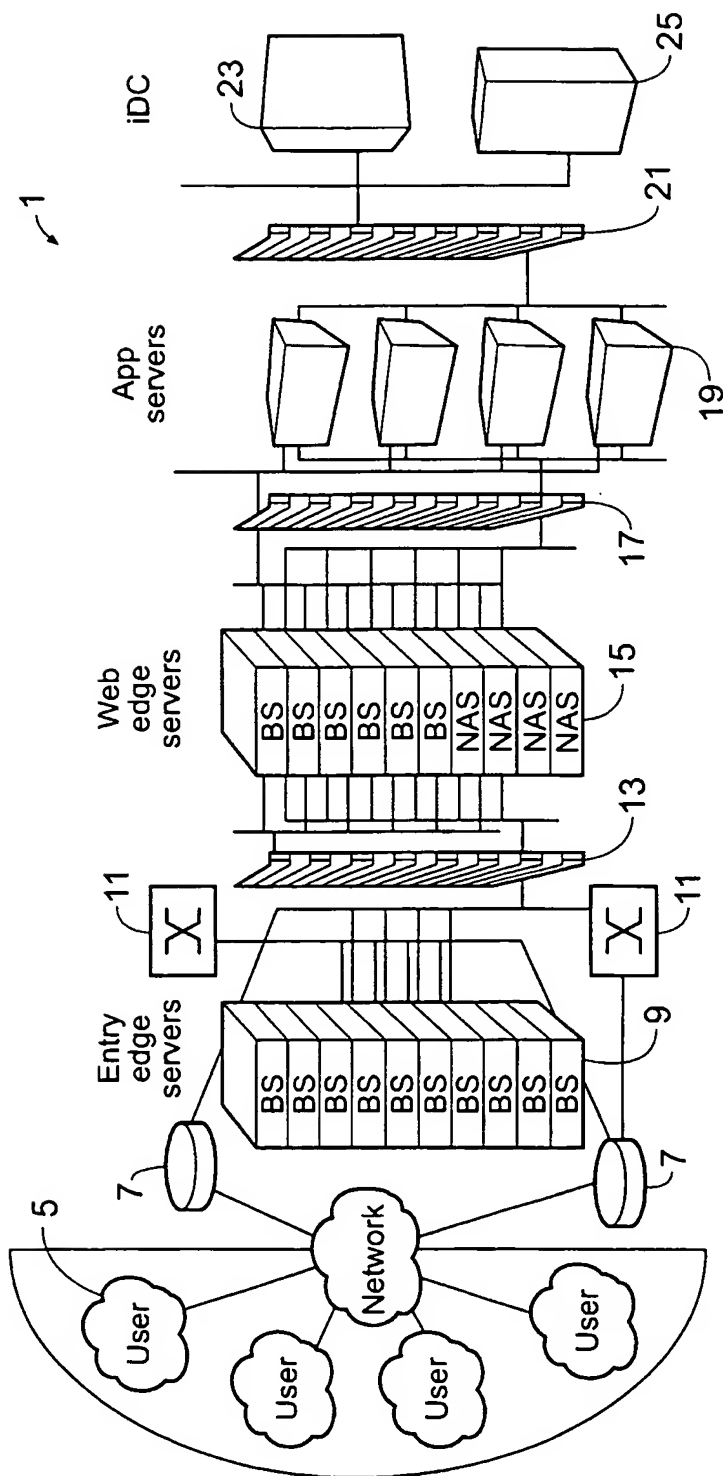


Fig. 1

3/25

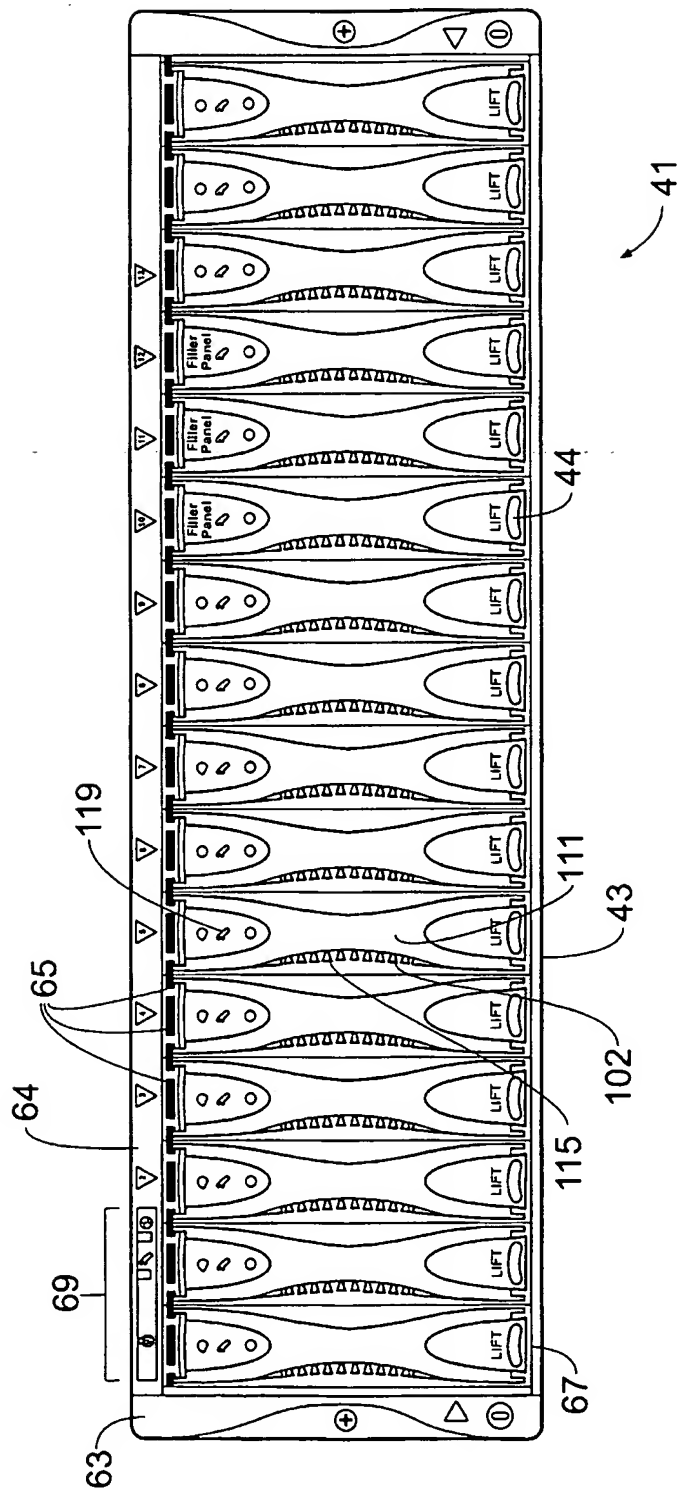


Fig. 3

4/25

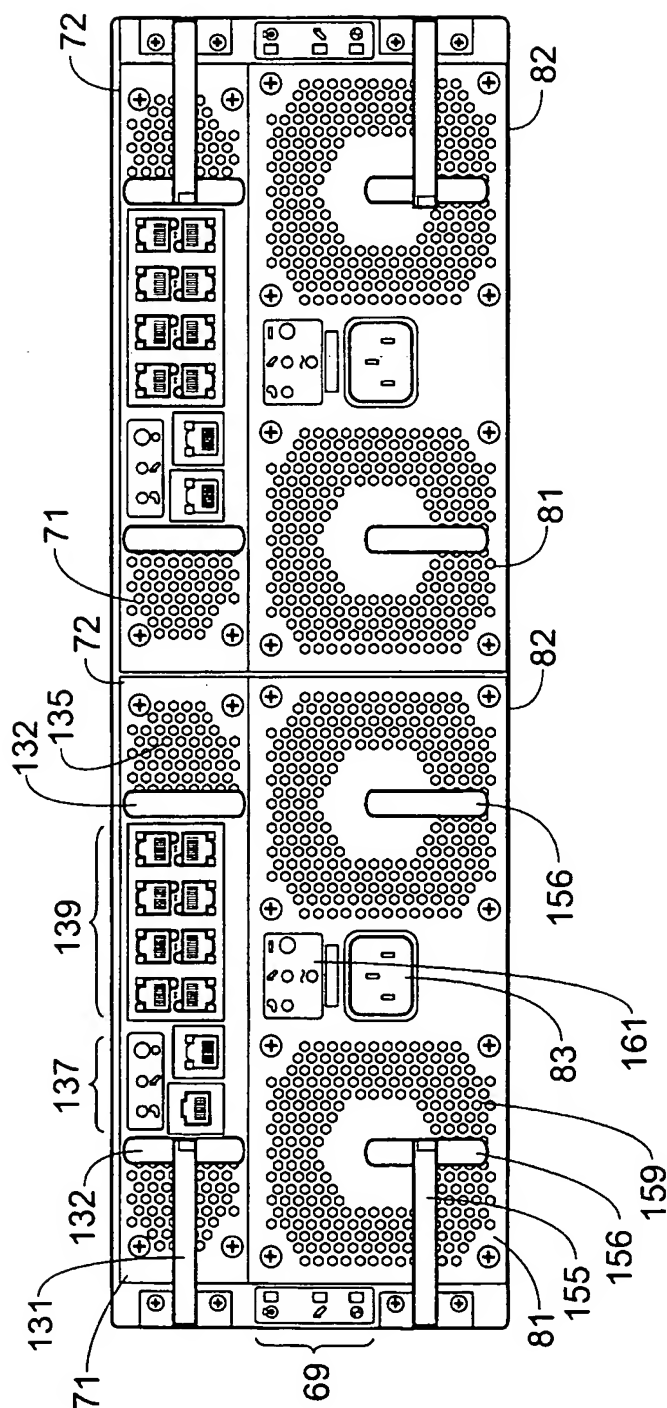


Fig. 4

5/25

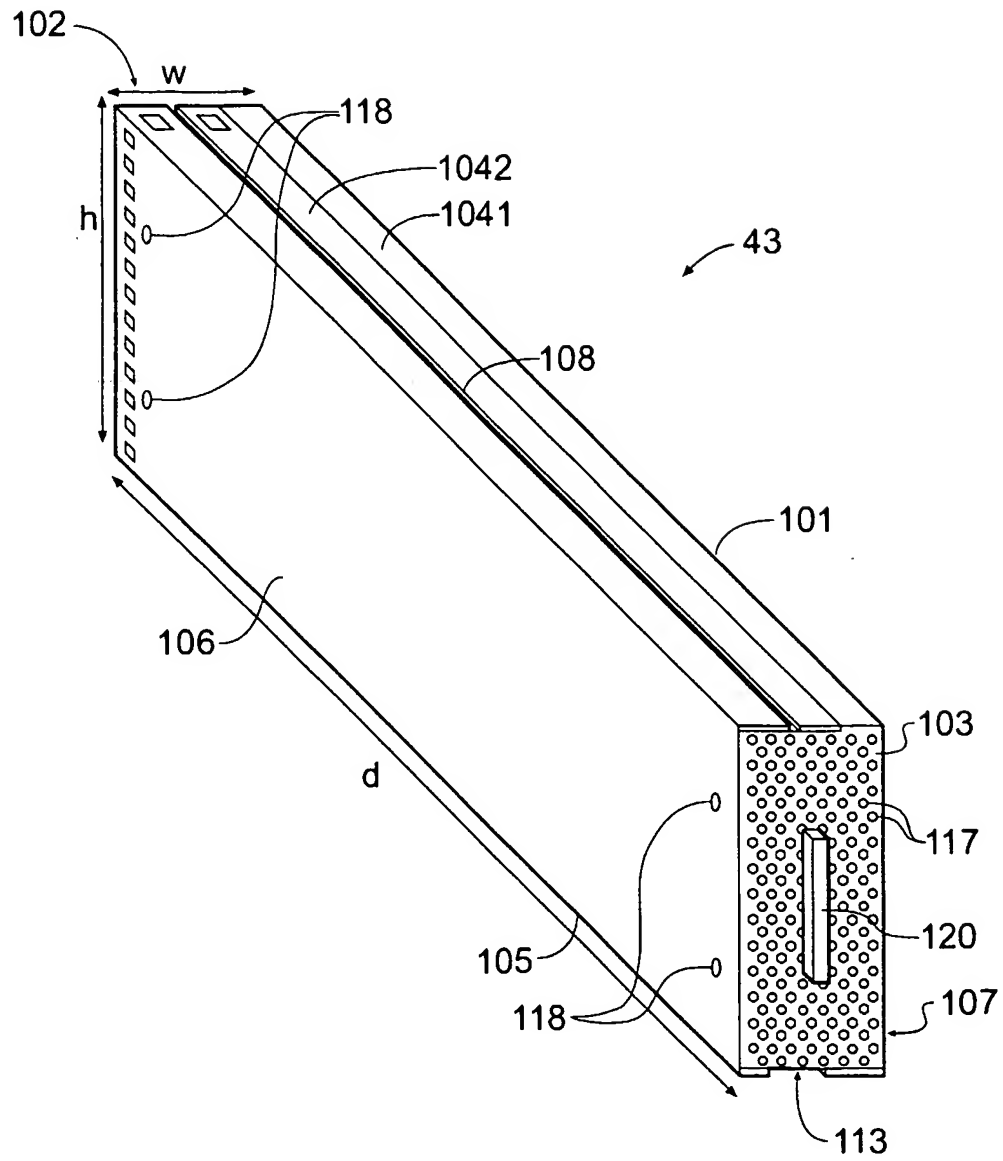
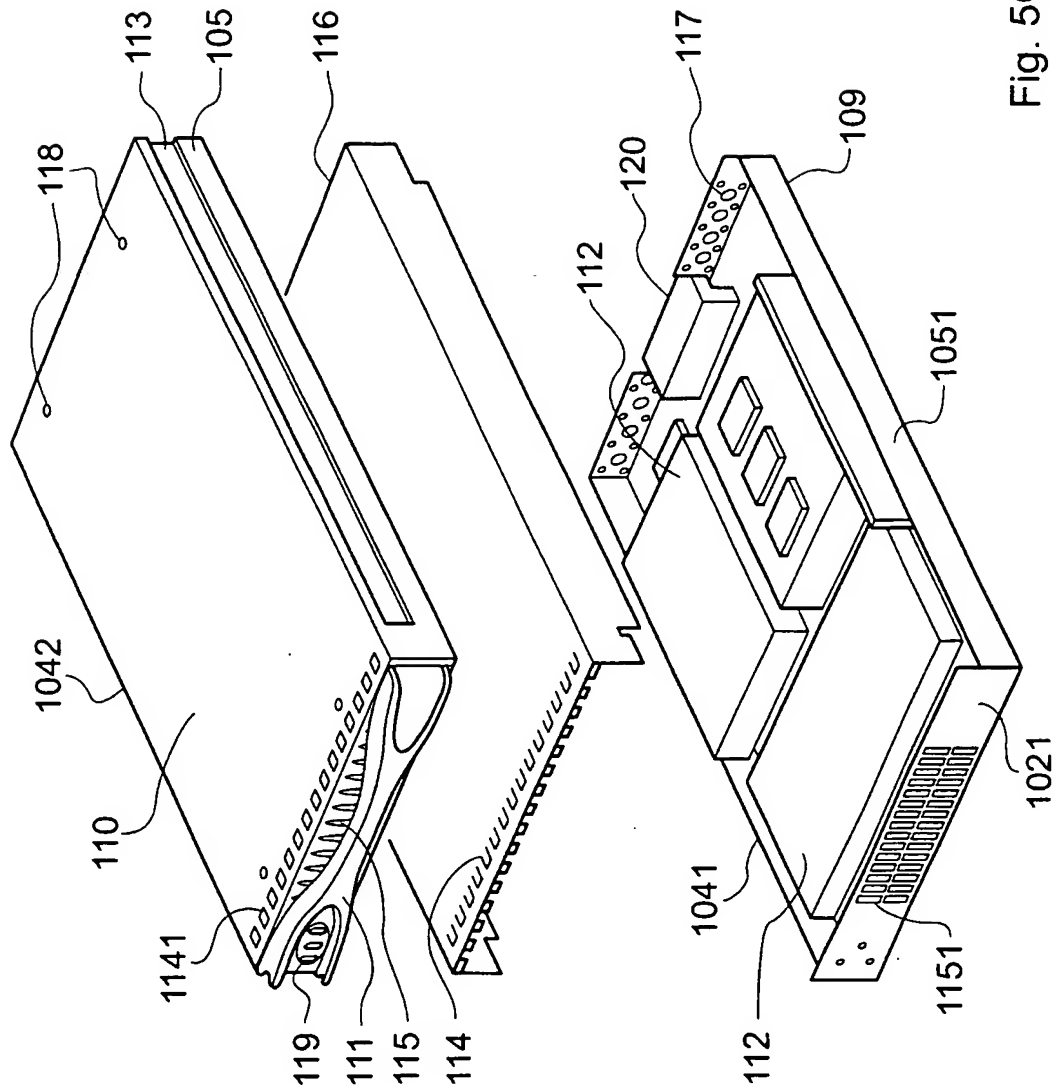


Fig. 5A

7/25



8/25

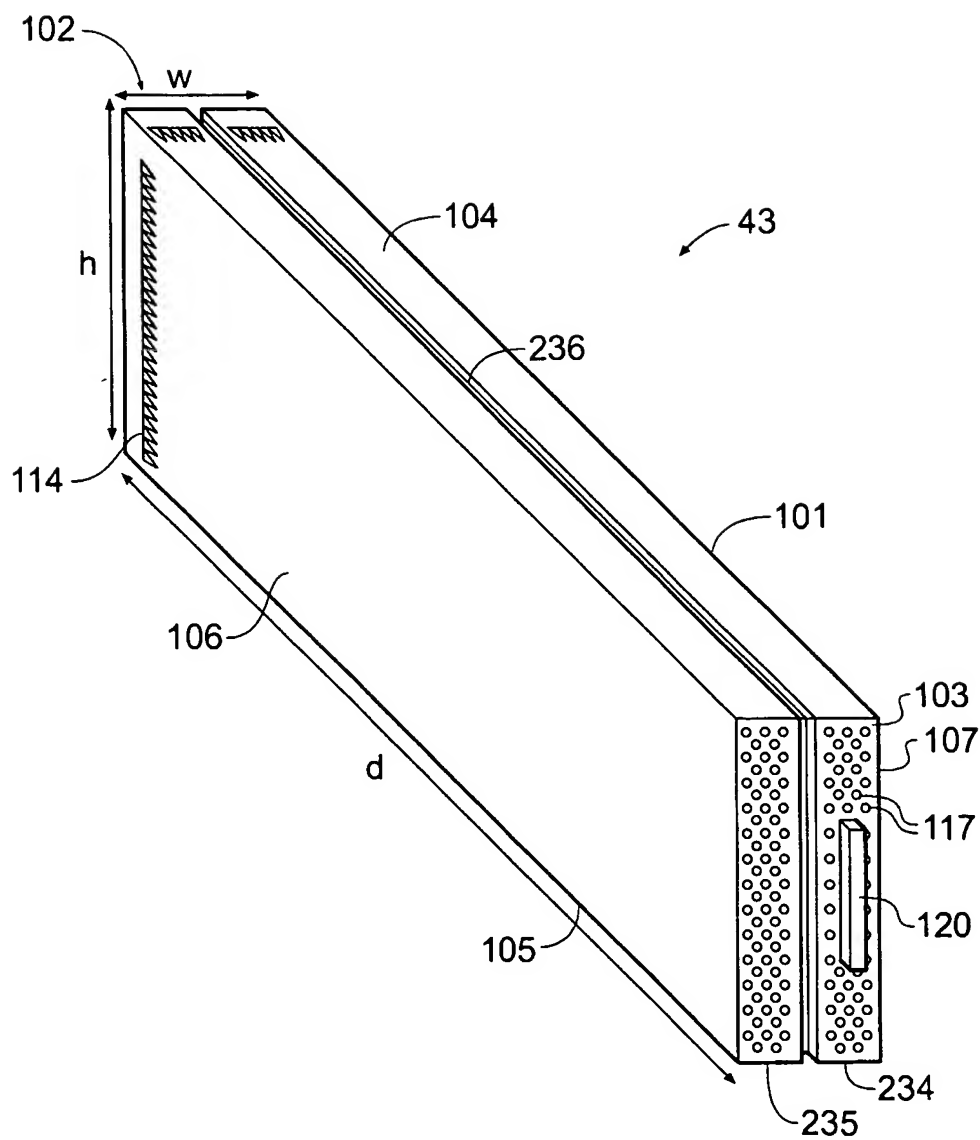


Fig. 5D

9/25

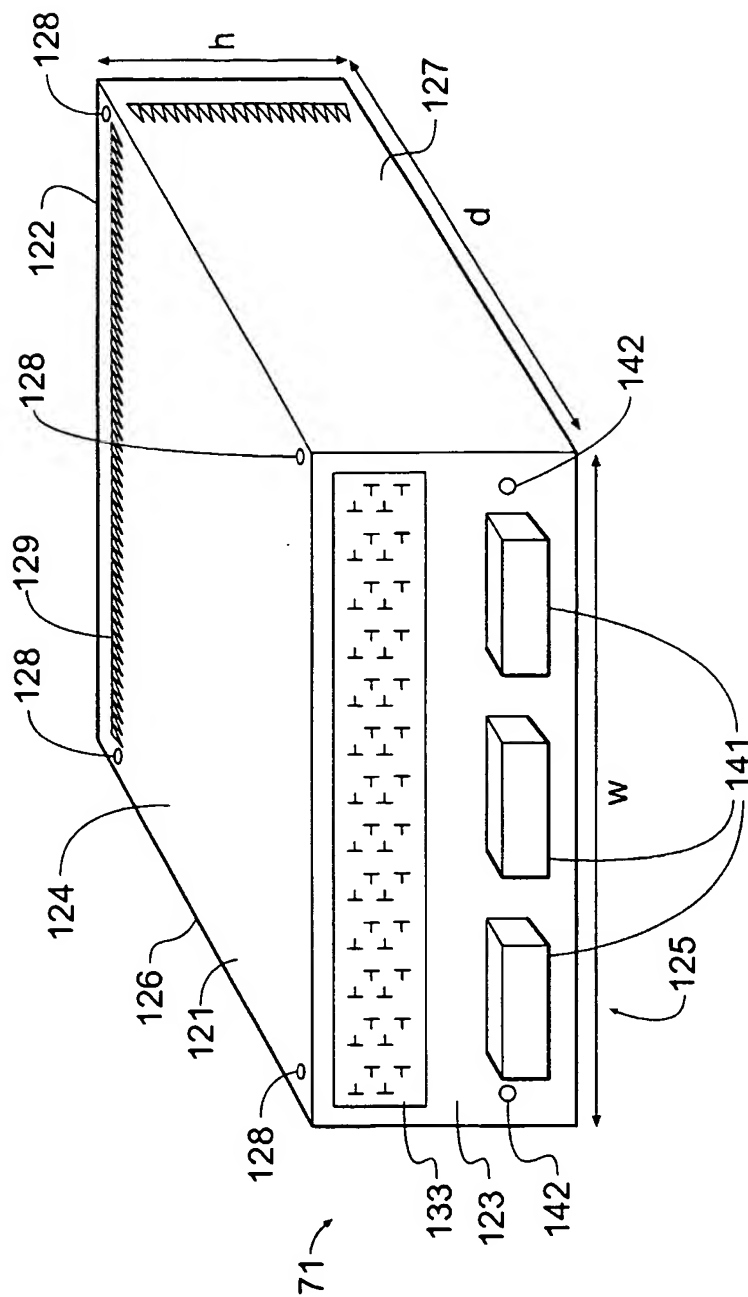


Fig. 6

11/25

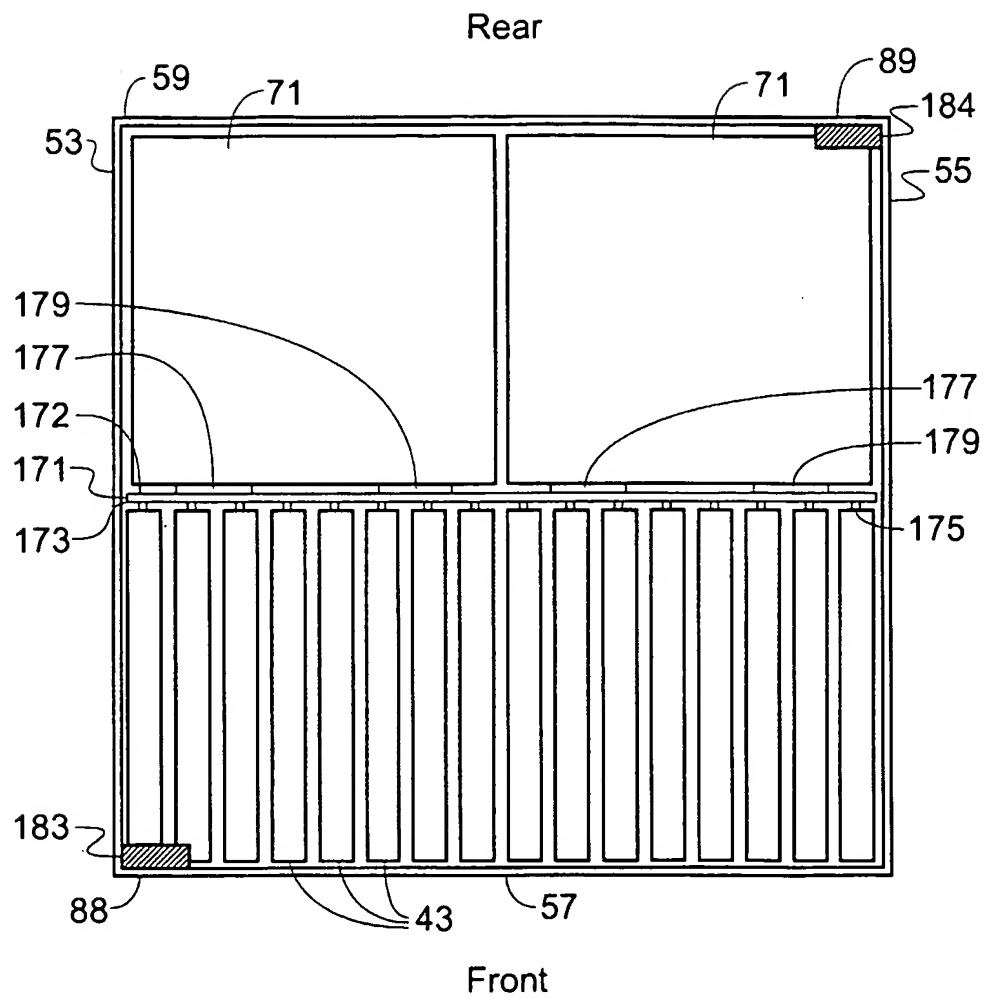


Fig. 8A

13/25

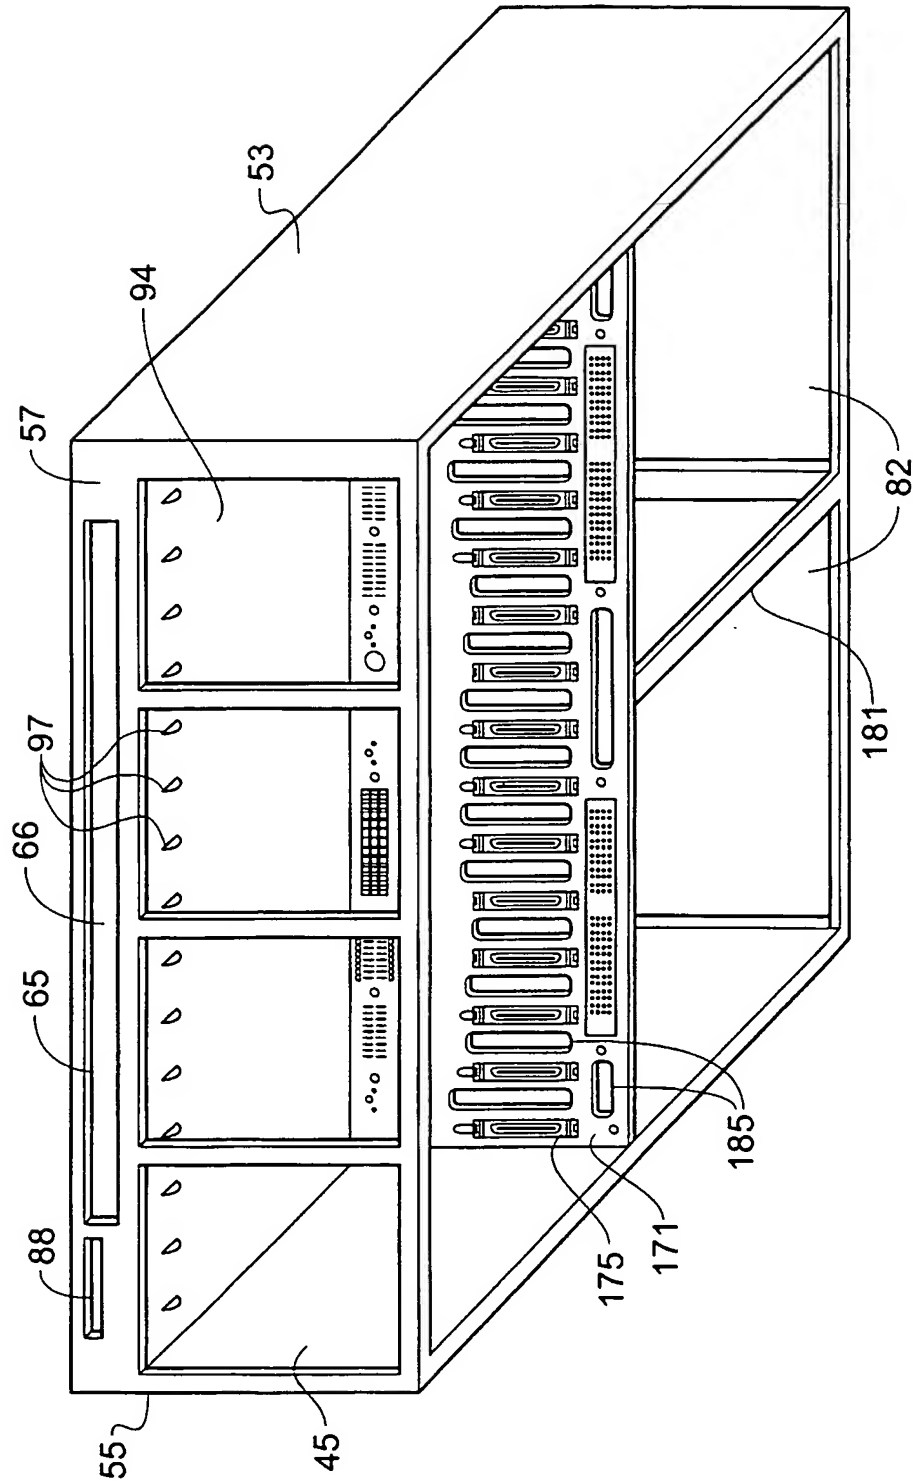


Fig. 8C

14/25

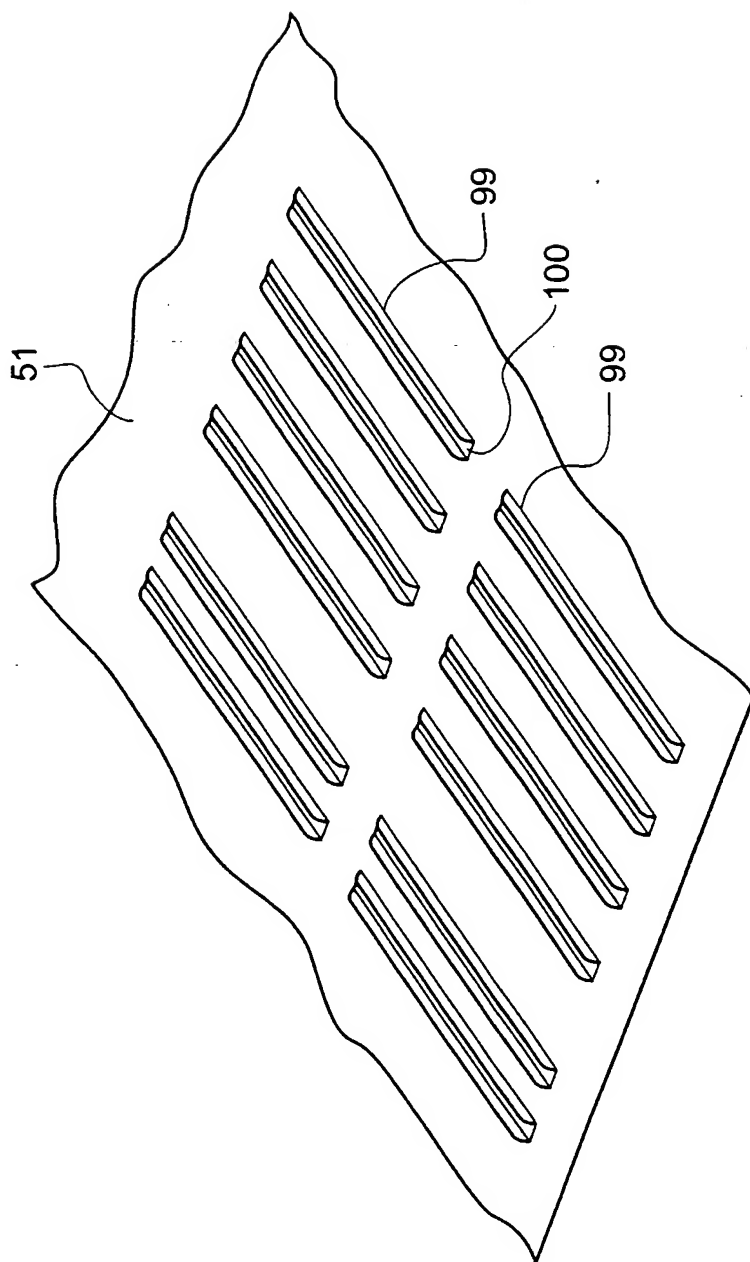


Fig. 8D

15/25

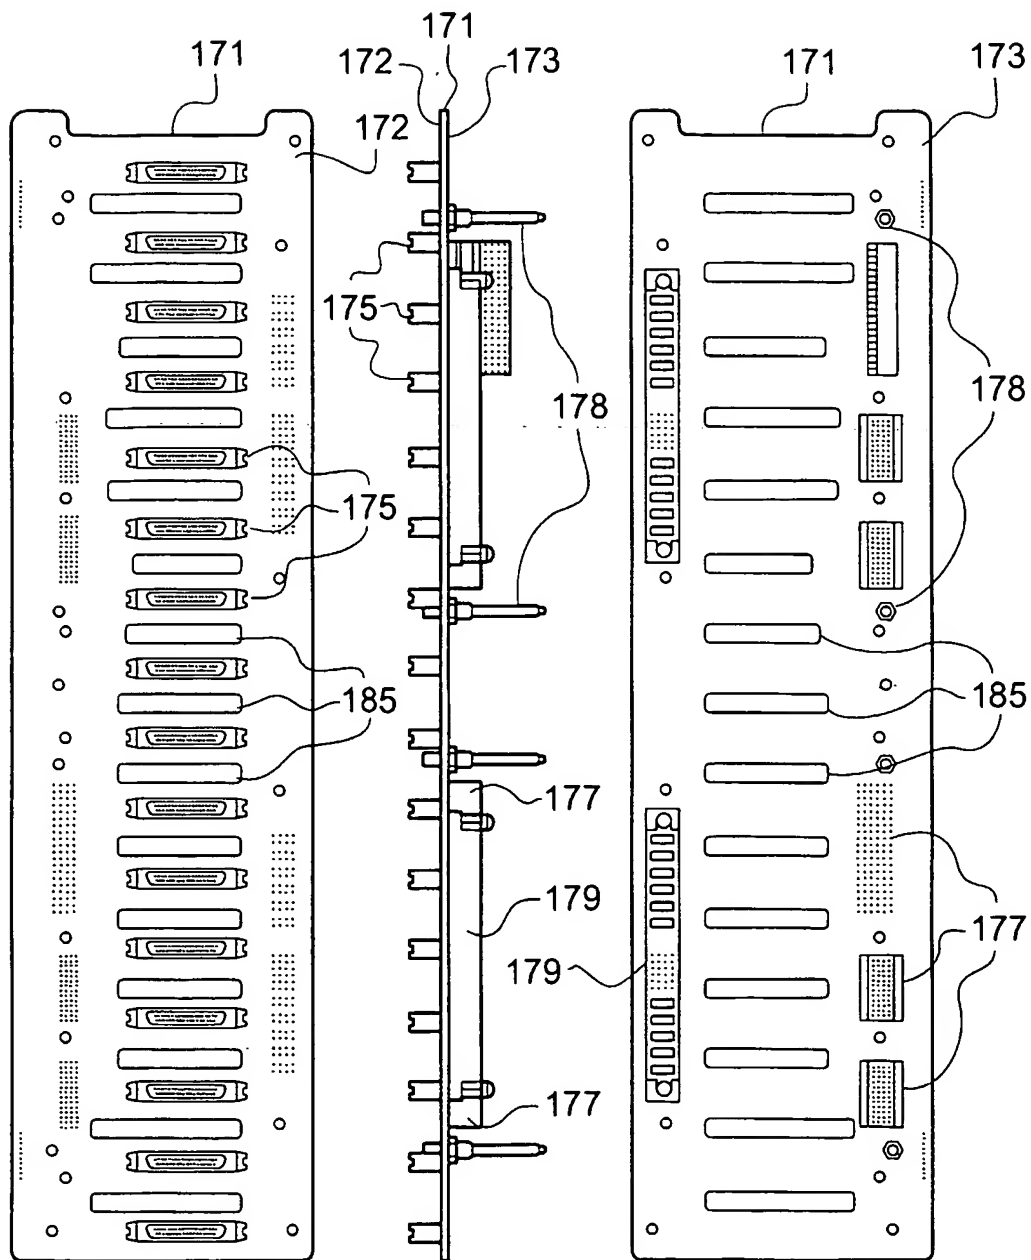


Fig. 9A

Fig. 9B

Fig. 9C

16/25

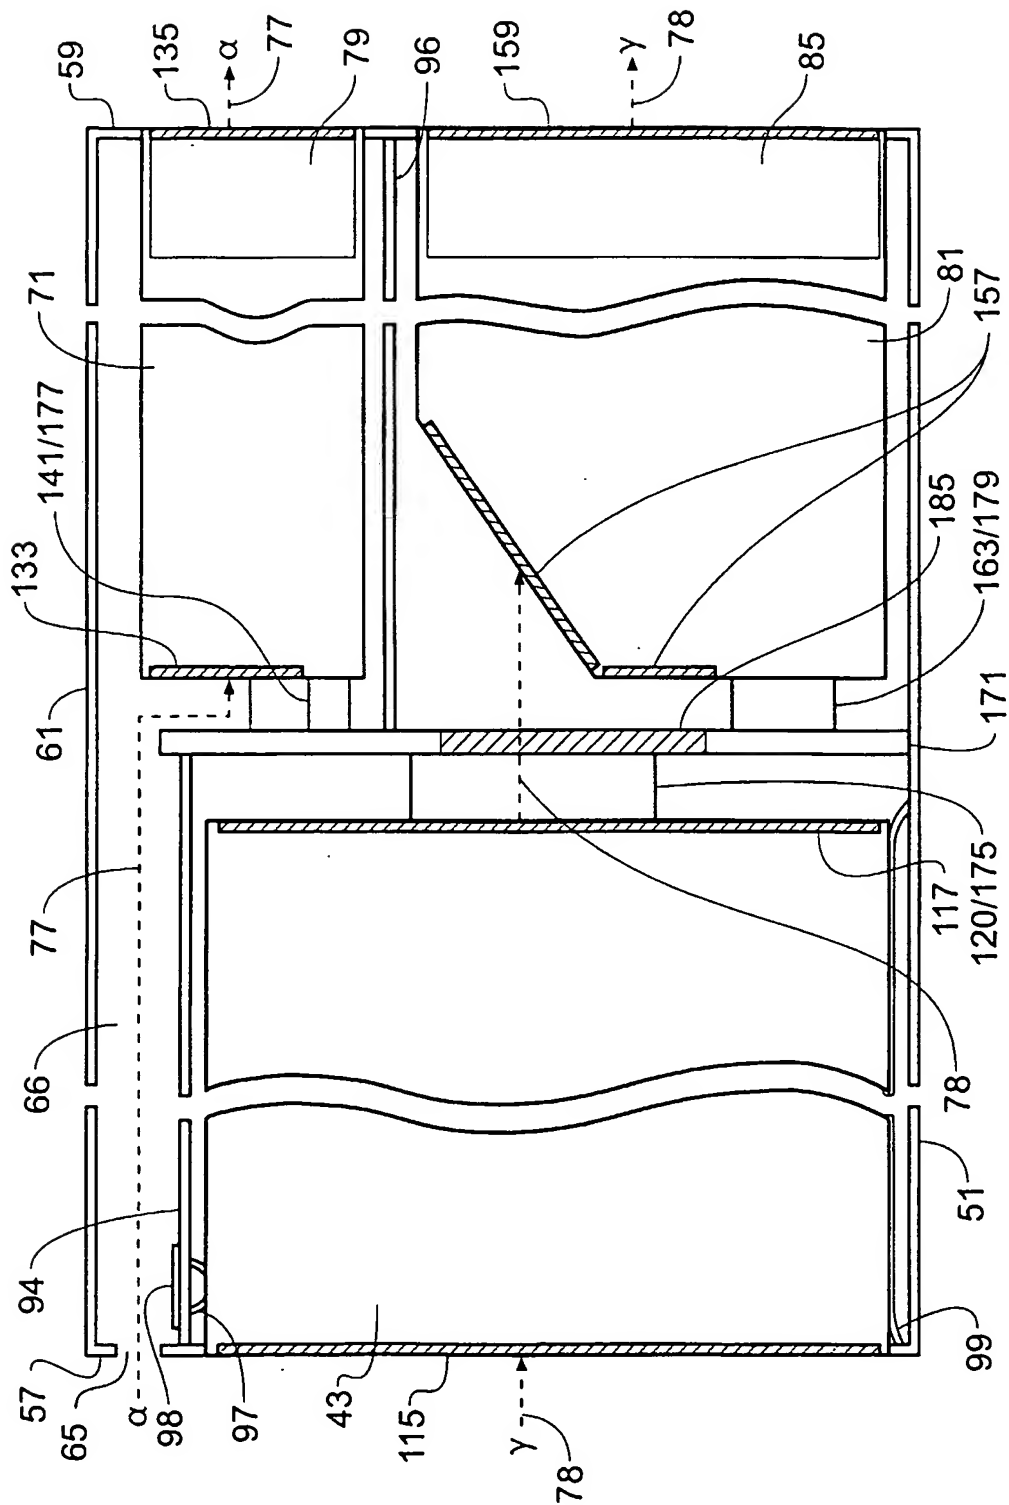


Fig. 10

17/25

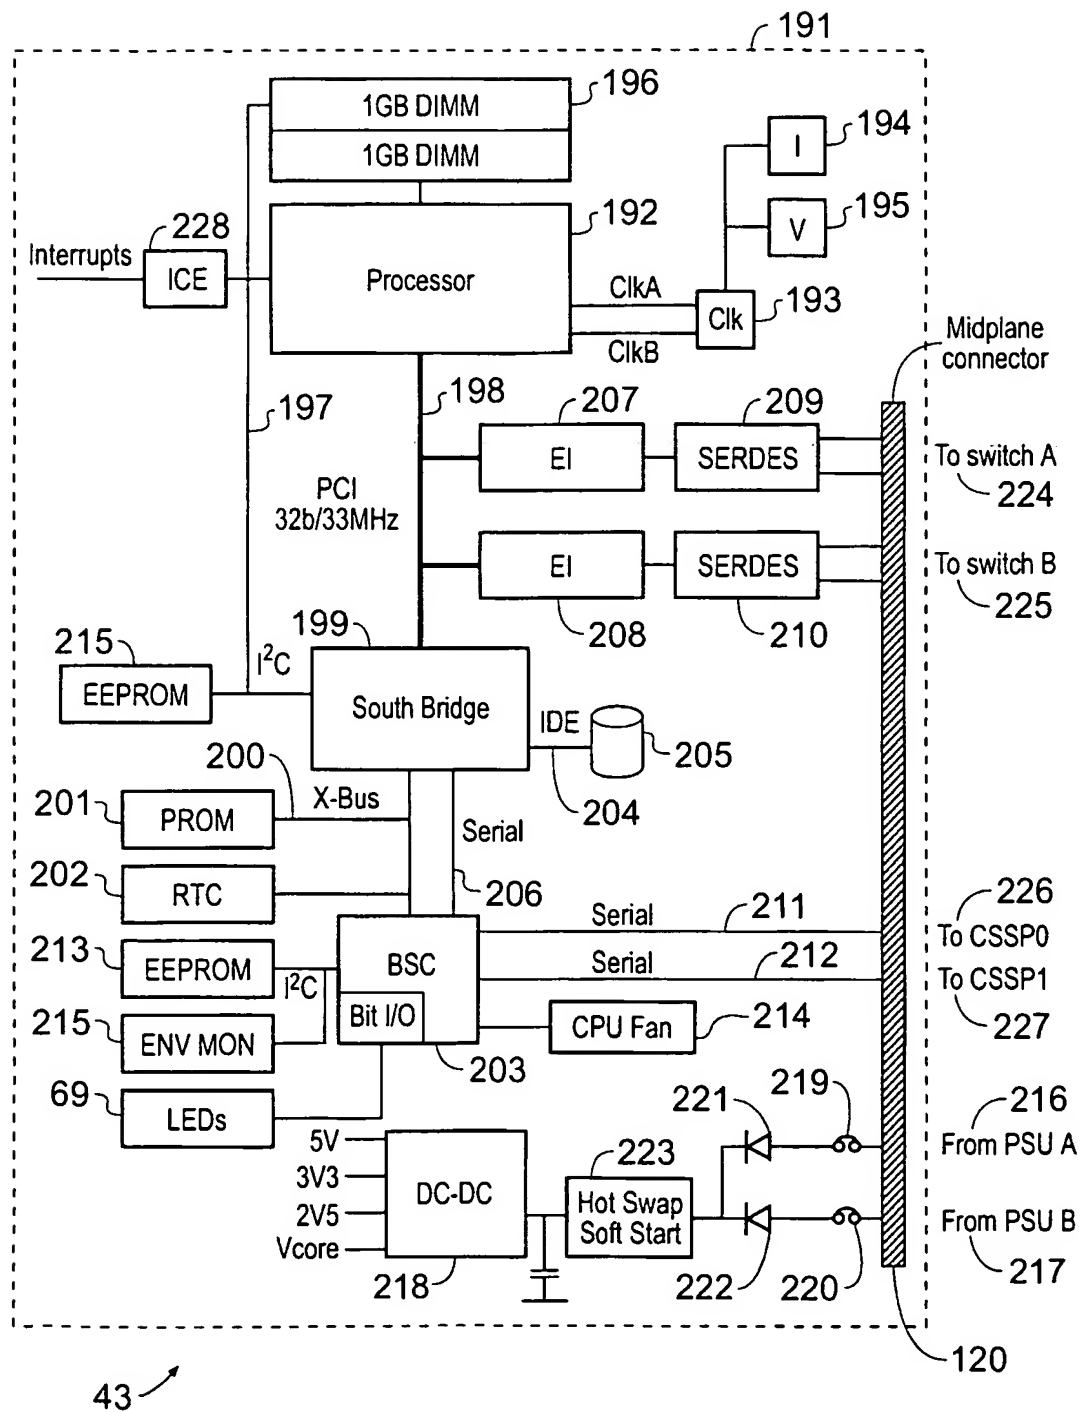
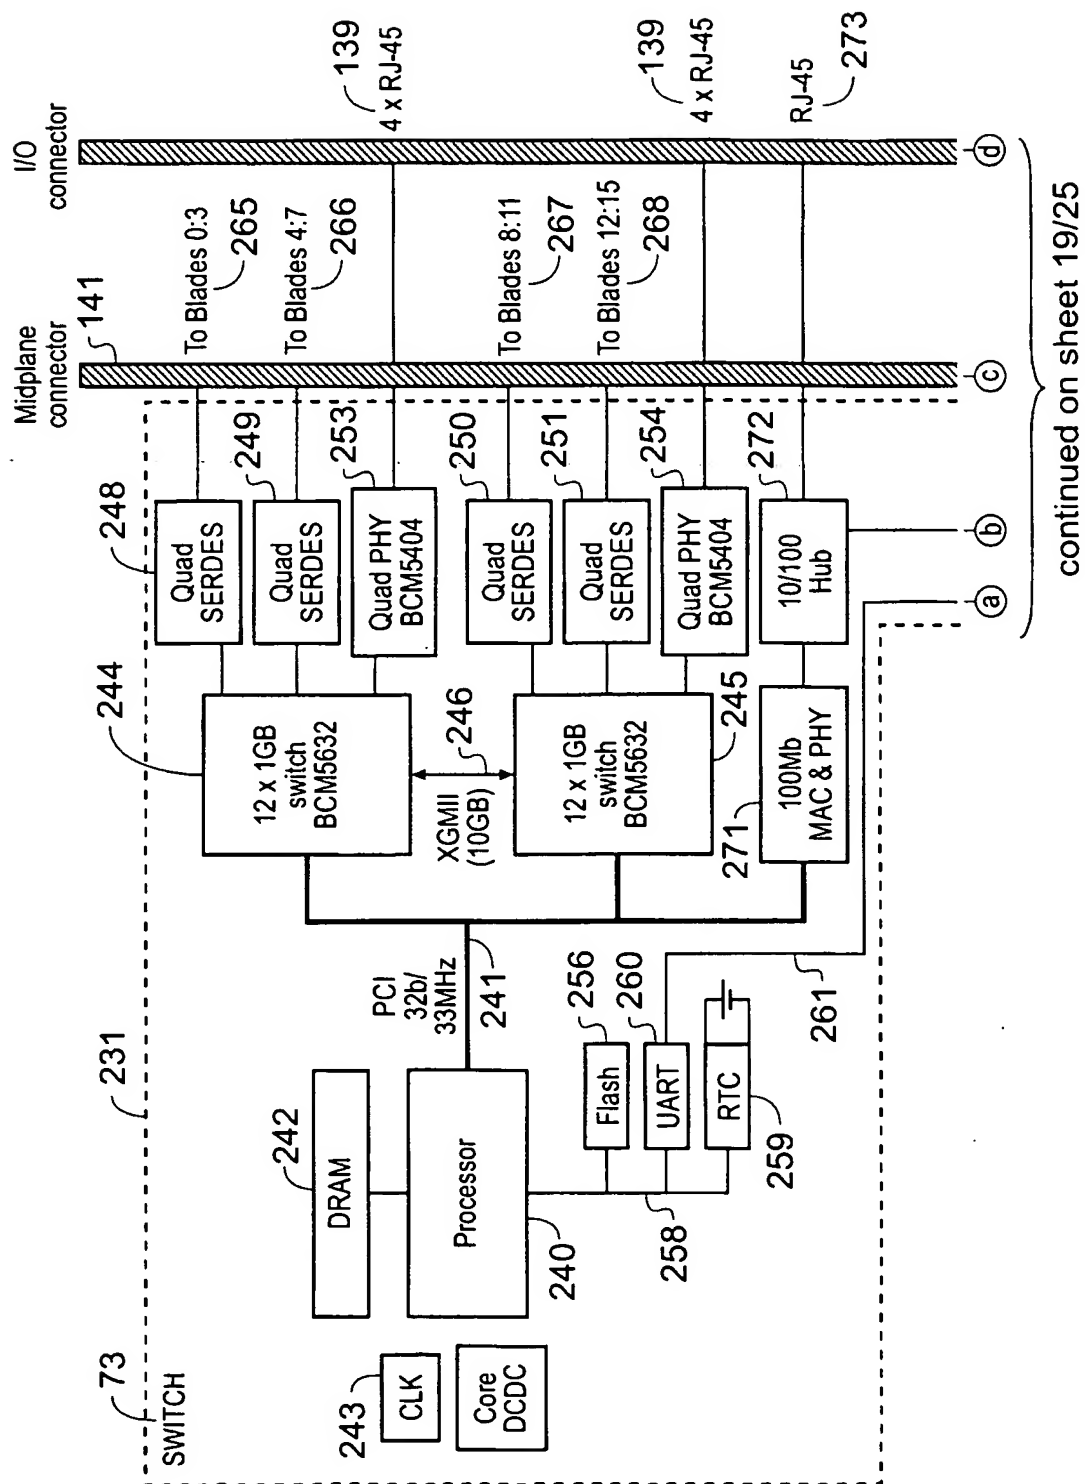


Fig. 11

18/25



19/25

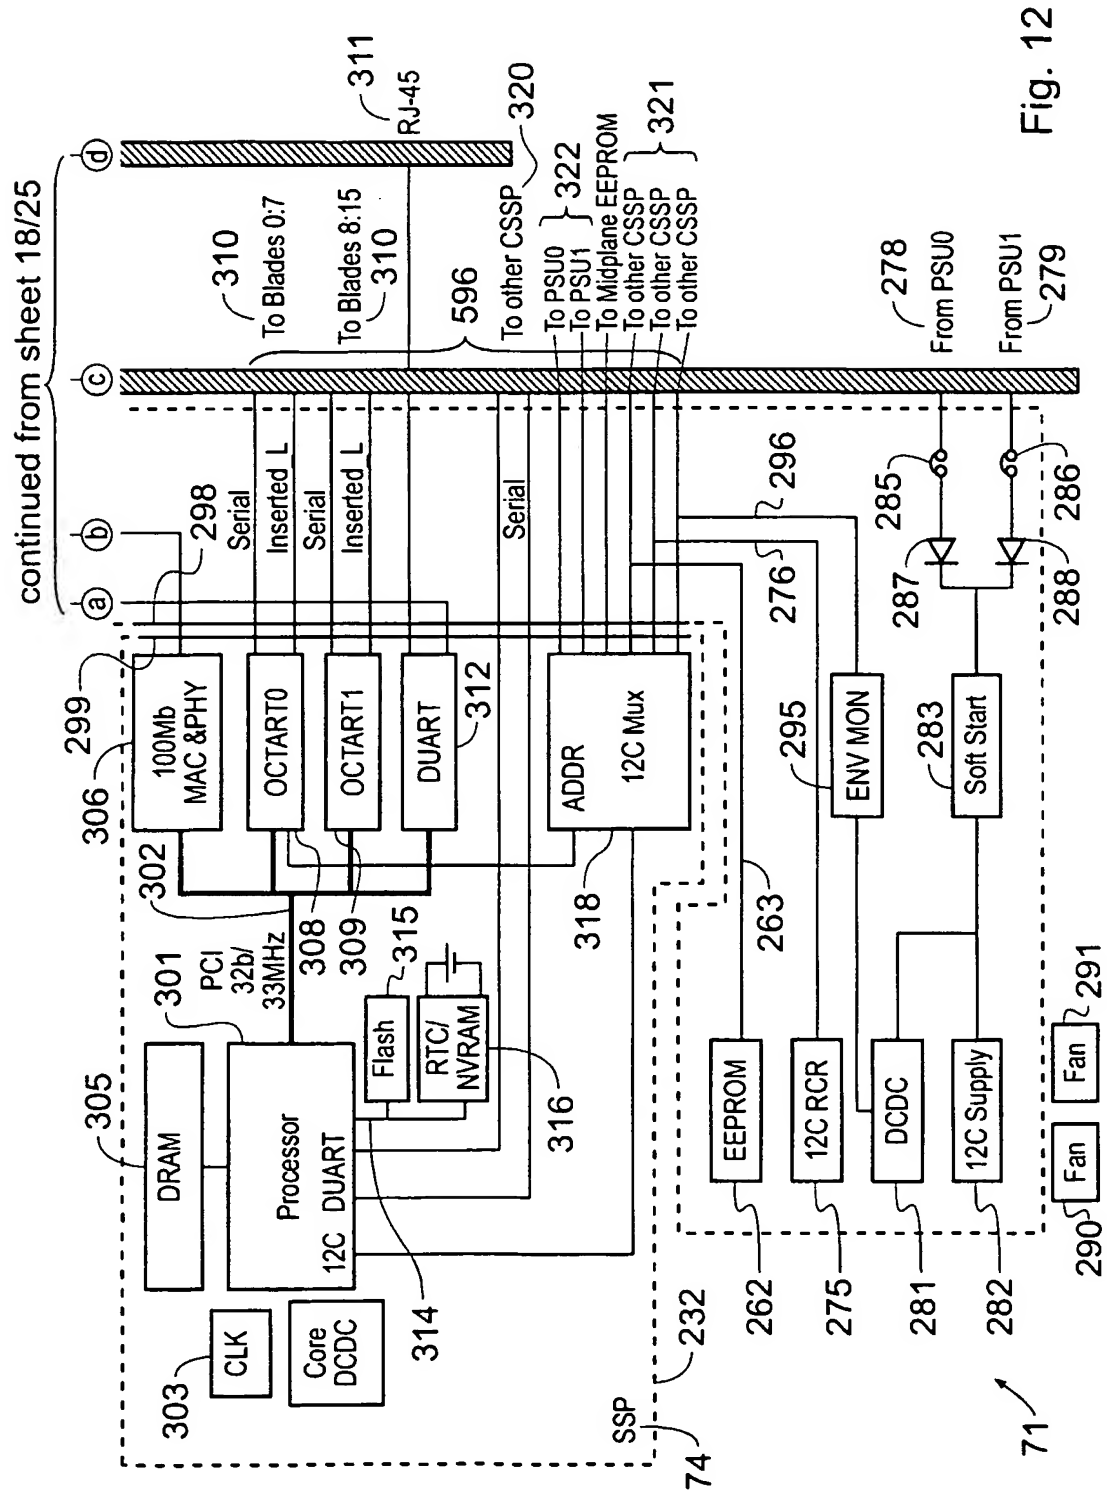


Fig. 12

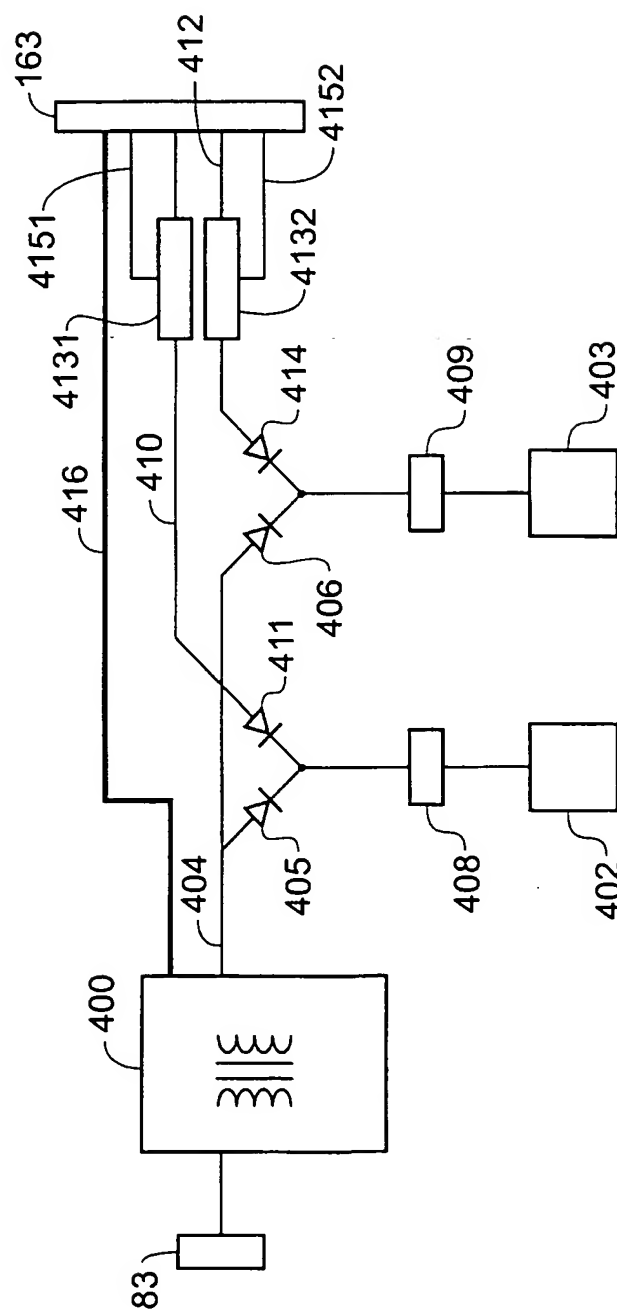


Fig. 13

21/25

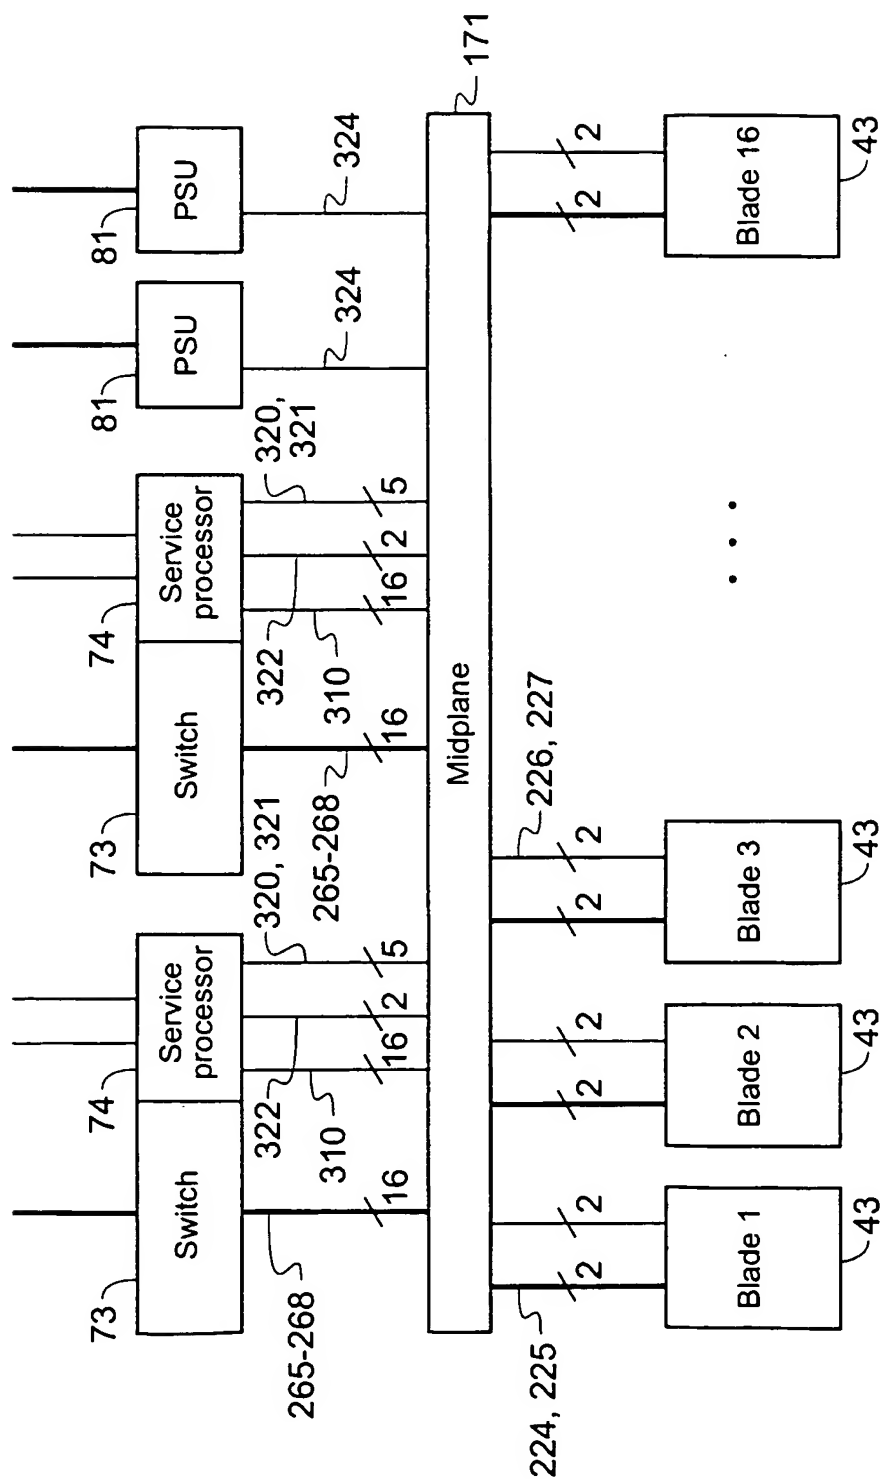


Fig. 14

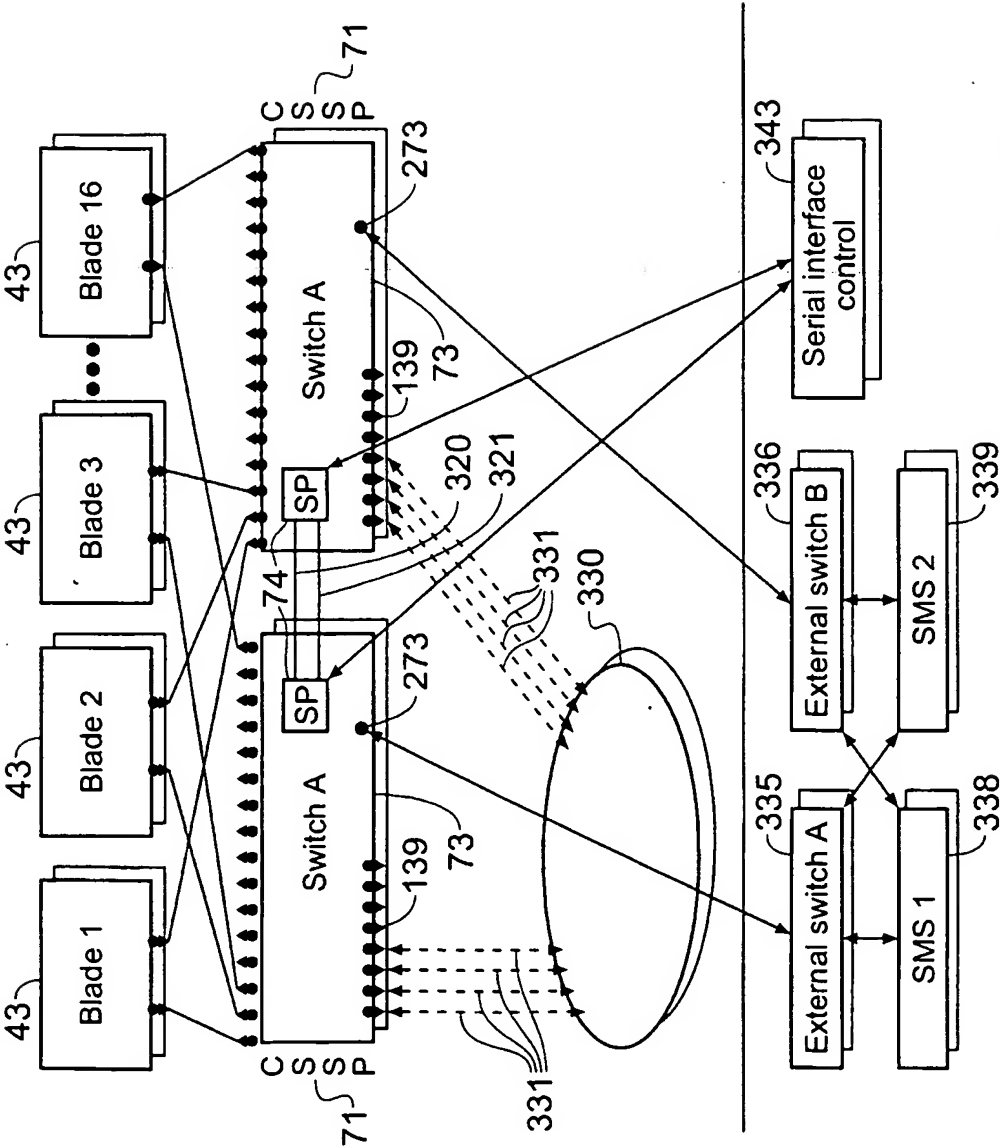


Fig. 15

23/25

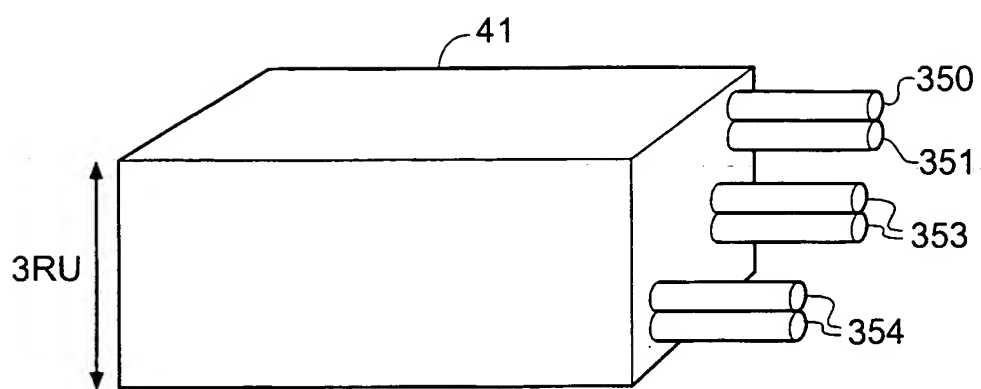


Fig. 16

24/25

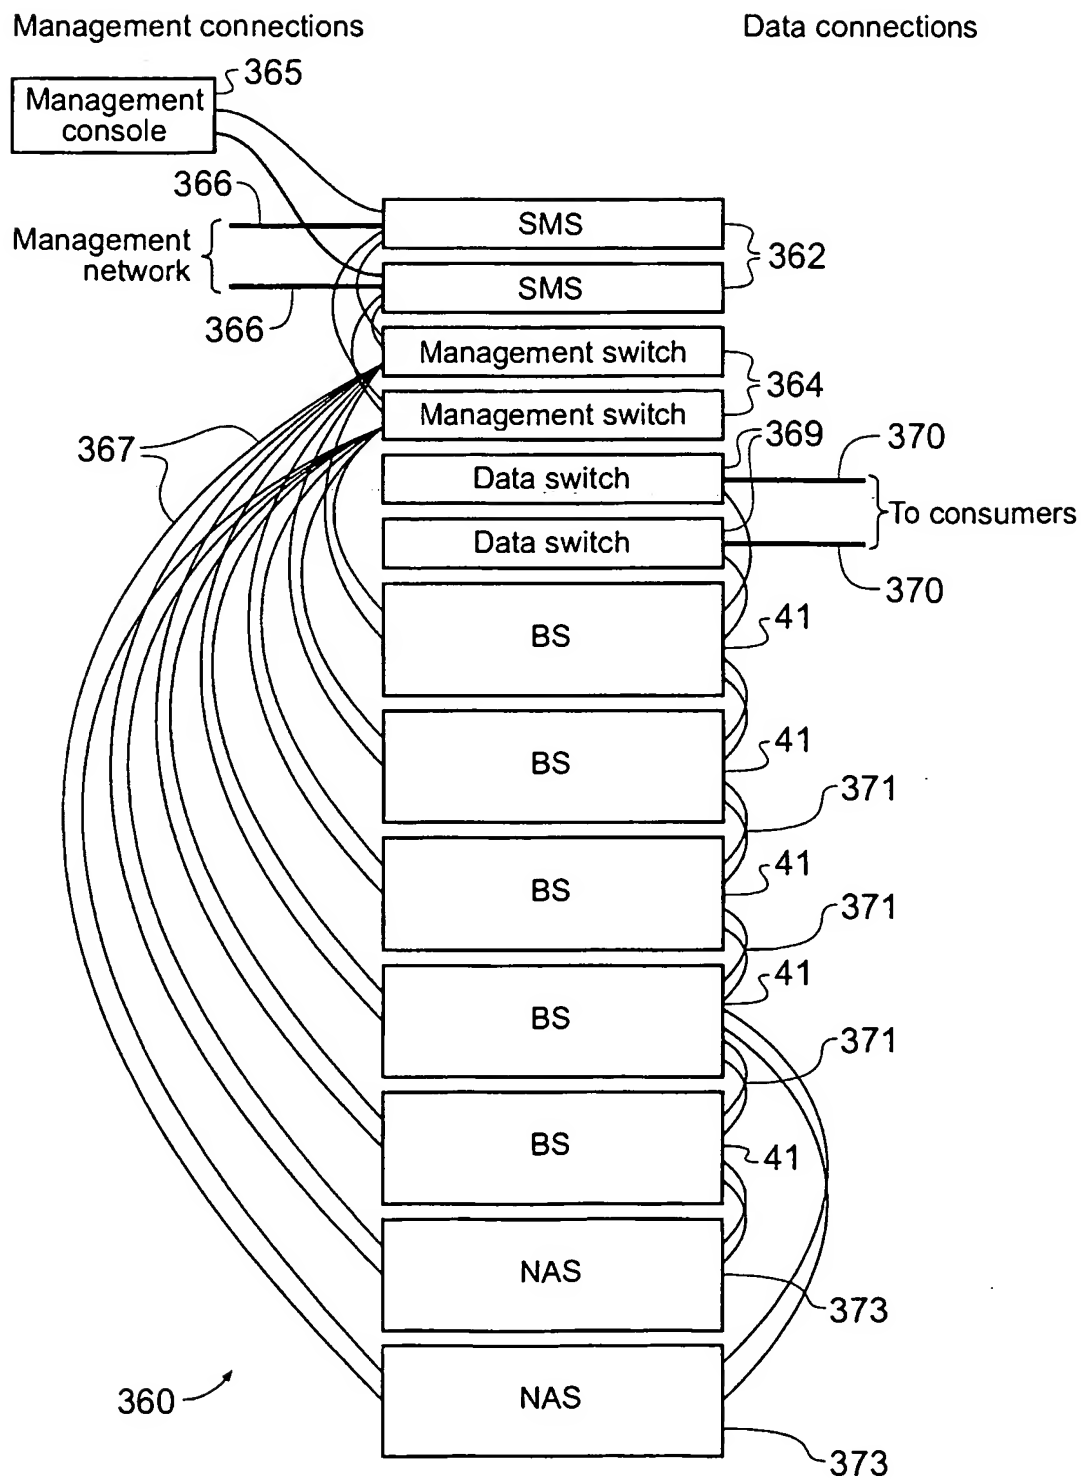


Fig. 17

25/25

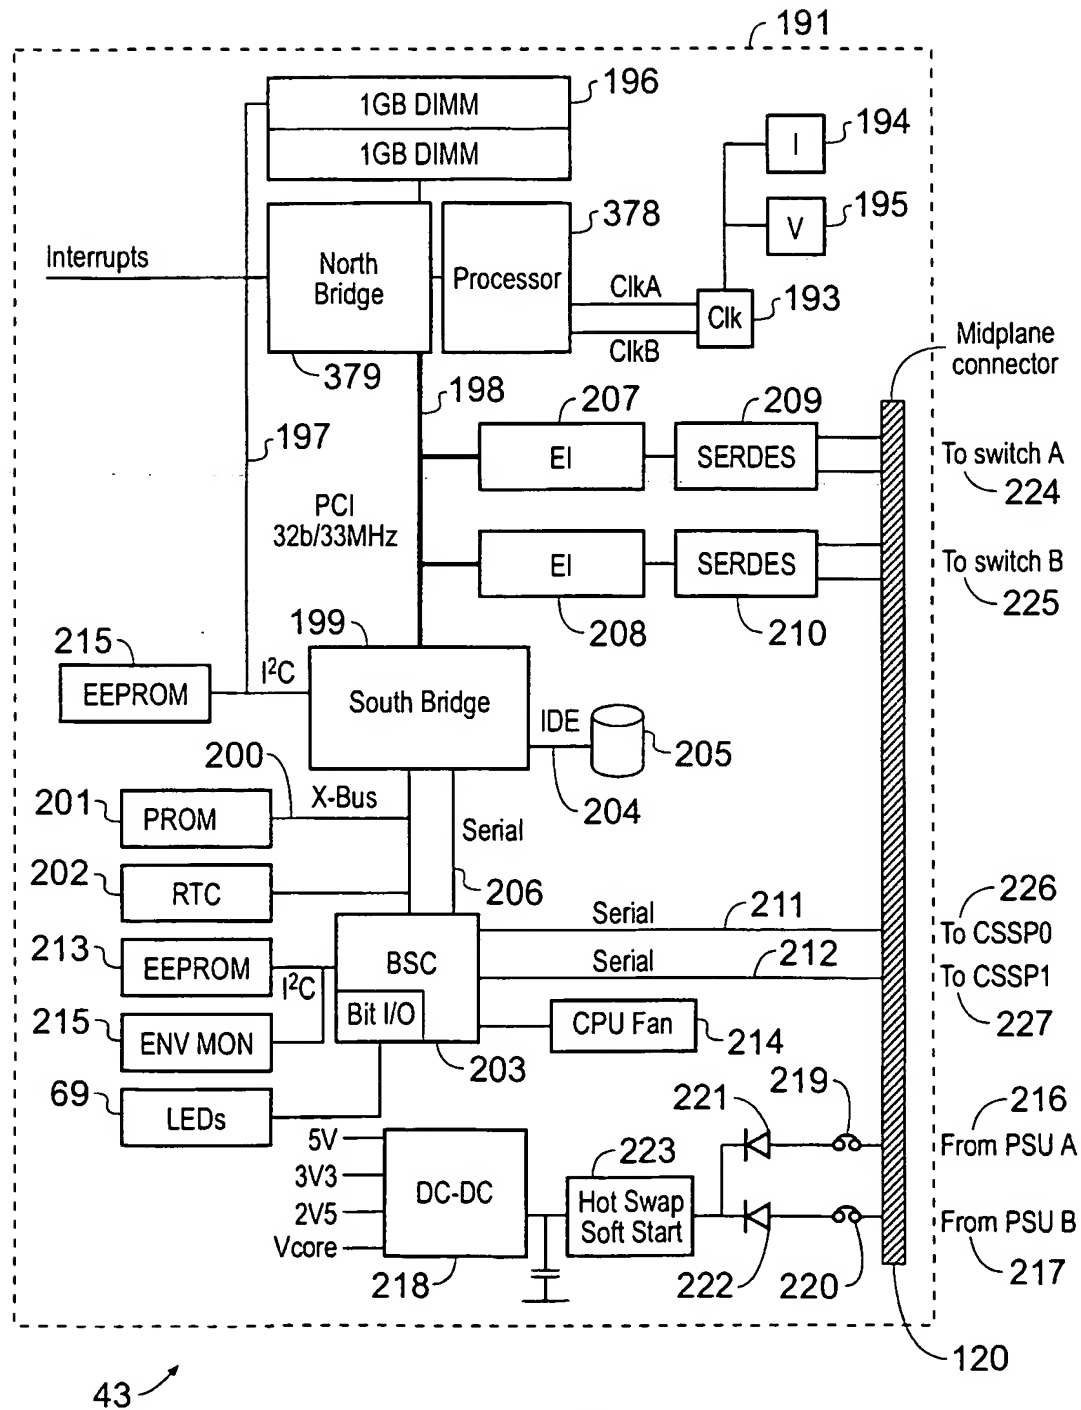


Fig. 18



SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

- (84) **Designated States (regional):** ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

SYSTEM MANAGEMENT

BACKGROUND OF THE INVENTION

5 The present invention relates to system management, and in particular but not exclusively to management of multi-processor systems, for example multi-processor server systems.

One application for the present invention relates to high density computer
10- systems, for example, computer server systems for telecommunications applications. In telecommunications applications, it is important to provide high reliability and high capacity of operation. Various approaches have been taken to providing such high-performance, high reliability systems. Typically such systems are designed around providing redundant resources so that if one component of the system develops a
15 fault, the system remains operational using the redundant resources. Fault tolerance can also be achieved, for example, with multiprocessor systems that provide redundancy through dynamic, e.g., software-controlled, task distribution. High density systems are typically rack mountable, with one or more processor systems occupying a shelf in the rack. The trend in recent times is to make the computers with
20 smaller form factors. This means that more computers can be located in a rack. This has the advantage of increasing the processing density within the racks, and also the advantage of reducing the distance between the computer systems.

The present invention relates to providing system management functions for a
25 modular computer system.

SUMMARY OF THE INVENTION

A first aspect of the present invention can provide a server system that includes a server blade carrier. The blade carrier can carry a plurality of removably mountable blades. Each blade can include a blade service controller operable to perform monitoring and management functions in the blade. The carrier can comprise at least one carrier service processor. The carrier service processor can be operably connected to the blade service controllers of blade servers mounted in the carrier to provide higher level management functions and can be configured to communicate management functions to and from the blades.

An embodiment of the invention can provide a hierarchical distribution of management functions within the server system, allowing management functions to be distributed to an appropriate level within the system. This management structure is readily scalable as the system grows. Also, it provides cost effective system management. The lower level management functions provided at a lower level have lower resource requirements to perform them than higher level functions provided at a higher level. It further enhances field replacement of blades and the insertion, in use, of new blades as a blade to be inserted include a service controller.

20

One embodiment provides the server system additionally comprising a system management server. The system management server can be operably connected to the carrier service processor of at least one blade carrier and the blade servers mounted in the carriers to provide high level management functions. This arrangement provides an additional tier to the management hierarchy allowing further distribution of management functions.

A further aspect of the present invention provides a method of providing management functions within a server system having a server blade carrier carrying a

plurality of removably mountable server blades, the method comprising: performing monitoring and management functions at a blade service controller within each blade; performing higher level management functions at at least one carrier service processor within the carrier; providing a communications channel connecting the carrier service
5 processor to each blade service controller; and communicating management functions between the carrier service processor and the blade service controllers.

Further aspects and advantages of the invention will become apparent from the following description of particular embodiments.

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments of the present invention will be described hereinafter, by way of example only, with reference to the accompanying drawings in which like reference
5 signs relate to like elements and in which:

Figure 1 is a schematic representation of an architecture of a multiprocessor system for supporting a web site;

10 Figure 2 is a schematic representation of a racking system incorporating an example of a carrier in the form of a rack-mountable shelf according to a first example;

Figure 3 is a front view of an example of a carrier in the form of a shelf of
15 Figure 2;

Figure 4 is a rear view of an example of the shelf of Figure 2;

Figures 5A, 5B and 5C are schematic perspective views and a schematic
20 exploded view respectively of an example of an information processing cartridge for mounting in the shelf of Figure 2;

Figure 5D is a schematic perspective view of an example of an information processing cartridge for mounting in the shelf of Figure 2;
25

Figure 6 is a schematic perspective view of an example of a combined switch and service processor module for mounting in the shelf of Figure 2;

Figure 7 is a schematic perspective view of an example of a power supply module for mounting in the shelf of Figure 2;

Figures 8A, 8B and 8C are a schematic plan view and schematic perspective views, respectively, of an example of the chassis and midplane of the shelf of Figure 2;

Figure 8D is a schematic perspective view of a part of a base member of the shelf of Figure 2;

Figures 9A, 9B and 9C are schematic front, top and rear views, respectively, of an example of a midplane of the shelf of Figure 2;

Figure 10 is a schematic cross section view through the shelf of Figure 2;

Figure 11 is a functional block diagram of an example of an information processing subsystem for the information processing cartridge of Figure 5;

Figure 12 is a functional block diagram of an example of an information processing subsystem for the combined switch and service processor module of Figure 6;

Figure 13 is a functional block diagram of an example of a subsystem for the power supply unit of Figure 7;

Figure 14 is a functional block diagram showing the connectivity between the components of the shelf of Figure 2;

Figure 15 is a functional block diagram showing the external connectivity of the shelf of Figure 2;

Figure 16 is a schematic representation of a shelf showing the external
5 connections from the shelf of Figure 2;

Figure 17 is a schematic representation of a rack mounted system comprising a plurality of such shelves; and

10 Figure 18 is a functional block diagram of a further example of an information processing subsystem for the information processing cartridge of Figure 5.

While the invention is susceptible to various modifications and alternative forms, specific embodiments are shown by way of example in the drawings and are
15 herein described in detail. It should be understood, however, that drawings and detailed description thereto are not intended to limit the invention to the particular form disclosed, but on the contrary, the invention is to cover all modifications, equivalents and alternatives falling within the spirit and scope of the present invention as defined by the appended claims.

DESCRIPTION OF PARTICULAR EMBODIMENTS

Embodiments and examples are described hereafter by way of example only in the following with reference to the accompanying drawings.

5

Shown in Figure 1 is an example of an application of a high capacity multiserver system 1 for implementing a network-connected web site such as, for example, an airline reservation system on the World Wide Web.

10 As shown in Figure 1, an external network 3 (e.g., the Internet) for communicating with a user 5 can be connected to gateways 7 which can be connected to an entry edge server group 9 implemented by a web farm as shown in Figure 1. The entry edge server group 9 forms an interface to the external network 3. The entry edge server group 9 can then be connected by switches 11 and a firewall 13 to a web
15 edge server group 15 that can also be implemented as a web farm as shown in Figure 1. The web edge server group 15 can serve to cache web pages that are readily accessible to users 5 accessing the system 1 from the external network 3, for example for checking flight times, etc. The web edge server group can comprise a number of blade server (BS) shelves and a number of network addressable storage (NAS) shelves
20 for storing critical data. The web edge server group 15 can be further connected by a further firewall 17 to a plurality of application servers 19, which can be responsible for, for example, processing flight reservations. The application servers 19 can then be connected via a further firewall 21 to computer systems 23, 25, for example, e-commerce services including financial services for receiving and processing payment
25 for airline reservations.

As will be appreciated, the server system described above with reference to Figure 1 is only an example of a possible application for a multiprocessor server system. Multiprocessor server systems have many different applications and the

present system is not limited to being applicable for use in only one or a limited number of such applications, rather multiprocessor server systems as described herein are operable for use in many different applications. A non-exhaustive list of such alternative applications includes: e-commerce web server systems; 5 telecommunications network server systems; LAN application and file server systems and remote vehicle control systems.

With reference to Figure 2, there is shown a schematic perspective representation of a rack system 31 as viewed from the front including left and right 10 front uprights 32 and 33 and left and right rear uprights 34 and 35. The uprights can be formed with apertures for receiving shelf fixings (e.g., screws, bolts, clips, etc., for mounting brackets, slides, rails, etc.).

Also shown in Figure 2 is an example of a blade server shelf 41 mounted in 15 the rack system 31. The shelf 41 forms a carrier configured to carry a plurality of information processing cartridges 43 located side by side along the shelf.

The term "shelf" is used herein in a conventional way to describe a structure that is mountable in rack system 31 and is configured to carry one or more 20 components to form at least a part of a rack-mountable system. In the present example, the shelf 41 is three-dimensional, having a height (H), width (W) and depth (D). In the present example, one dimension (hereinafter described as the height, H) is smaller than the other dimensions (hereinafter described as the depth, D, and the width, W) to facilitate mounting of the shelf within the rack system 31. It will be 25 appreciated that although the width and depth are typically constrained by the dimensions of the racking system for which the shelf is designed, there is more freedom as regard the height, subject to taking account of any appropriate standards and packaging considerations.

Each of the information processing cartridges contains at least one processor. Each information processing cartridge in the present example is operable as a server. In the described examples, the information processing cartridges are configured as robust enclosed modules.

5

In the example to be described in more detail presently, the information processing cartridges, when aligned in the carrier shelf, look like rectangular slabs, or blades. Accordingly, an information processing cartridge can be described as a blade.

The information processing cartridges 43 comprise information processing modules
10 enclosed in an enclosure, or housing, so that the information processing modules have the form of cartridges. Also, as the information processing cartridges are to operate as computer servers in the example described in more detail presently, an information processing cartridge 43 can also be described as a server blade. Accordingly, in the context of this example, the terms module, cartridge and blade are used
15 interchangeably.

The illustrated example of a shelf 41 is configured to carry sixteen information processing cartridges 43, each of which is removably mountable in a respective opening 45 in the front of the shelf, whereby the information processing cartridges can
20 be inserted into and removed from the front of the shelf 41 without removing the shelf 41 from the rack system 31.

In the present example, the shelf 41 comprises a three-dimensional, generally rectangular, enclosure, or housing, 47 that is suitable for mounting in generic racking
25 systems including both 4-post and 2-post systems. It can be mounted on fixed rigid rack mounting ears and/or a simple slide/support system. The present example is designed for standard 19"-wide racking (1"=25.4mm) as defined, for example, in the well-known IEC297 and EIA 310 specification standards with height corresponding to the so-called 3U (3 standard unit) height. For mounting such a 3U unit in such a 19"-

wide racking system, with a depth of, say 25" or 30", the enclosure can be arranged with a height of up to about 130.5mm, a width of up to about 445 mm and a depth, including all hardware and fascias, but excluding cable management, of up to about 635mm, with the depth from the front-most point of a fascia to a rear I/O connector
5 panel of a rear mounted Field Replaceable Unit (FRU) of about 610mm. Of course, other examples designed for other racking systems could have different dimensions.

This example of a shelf 41 has a single enclosure, or housing, 47 that houses a number of modular units or subsystems, the majority of which are replaceable in the
10 field and are therefore known as Field Replaceable Units (FRUs). These modular units include the information processing cartridges 43.

The shelf enclosure 47 can be fabricated from sheet material (e.g., from steel sheet) to form a chassis portion 49 that includes a base 51, two sides 53 and 55, a
15 front 57 and a rear 59. The word "front" as used here is merely used as a label herein to refer to the face, or wall 57 of the enclosure that is located at the main access side of the rack system 31 in use when the shelf is mounted therein. Similarly, the words "rear" and "side" are merely used as labels herein to refer to the faces, or walls 59, 53 and 55 that, in use, are located at those respective positions when the shelf is mounted
20 in the rack system 31.

The openings 45 can be formed in the front face 57 for receiving the information processing cartridges 43 and, as will be explained later, apertures can also be formed in the rear face 59 for receiving further FRUs. The enclosure can further
25 include a removable top cover 61 that can be secured to the chassis portion 49 by suitable fastening (e.g., screws). The apertures in the front and rear faces 57 and 59 allow at least some of the FRUs to be inserted into and/or removed from the shelf enclosure 47 via the front or the rear thereof, as appropriate, without removing the shelf from the racking. Access to components mounted in the shelf that are not

accessible via one of the apertures in the front 47 and rear 59 faces can be achieved by removing the shelf enclosure 47 from the racking system 31 and then removing the top cover 61 of the shelf enclosure 47.

5 Figure 3 is a front view of an example of a shelf 41 for a first example. A plastic front bezel 63 can be provided that fits on the front face 57 (shown in Figure 2) of the chassis 49 of the shelf enclosure 47. The front bezel 63 can be formed as a unitary removable part that spans the whole width and height of the front of the shelf enclosure 47. The front bezel 63 could alternatively include a number of separate
10 components, or mouldings. The front bezel can include a peripheral portion 64 that can provide areas for corporate and product branding marks, for identification and numbering for the information processing cartridge and for a bar code label (all not shown). One or more apertures 65 can be formed in the peripheral portion 64 of the bezel 63. The apertures 65 in the bezel can be arranged to align with one or more
15 apertures (e.g. a slot (not shown in Figure 3) in the front face of the chassis. In use, air can pass through the apertures 65 to flow into the shelf enclosure 47 to reach FRUs that are mounted in the shelf enclosure 47 through the rear face 59 thereof. Air flowing through the aperture 65 flows into a plenum chamber 66 (not shown in Figure 3) to flow past the processing cartridges 43 to reach rear mounted FRUs. A central
20 area 67 of the front bezel 63 can be open allowing access to the openings 45 in the front face 57 of the shelf enclosure 47 for insertion and removal of the information processing cartridges 43. Where no active module is mounted in a location for an information processing module, a blanking panel, or filler panel, such as the blanking panel 44, can be located in that location. LED indicators 69 can be mounted on a
25 system indicator printed circuit board (not shown) behind a designated area of the bezel to provide an indication of system status via light guides incorporated into the bezel. A further system indicator board (also not shown) carrying LED indicators can be provided inside the shelf enclosure to be visible from the rear thereof.

As mentioned above, in the present example of a shelf, up to sixteen information processing cartridges 43 can be installed in respective openings 45 in the front face 57 thereof. The number of information processing cartridges 43 actually installed in any installation is dependent upon the system configuration required.

5 Various features relating to the information processing cartridges 43 that are shown in Figure 3 will be described later.

Figure 4 illustrates the rear of the shelf unit of Figures 2 and 3. This shows two different types of FRU 71 and 81 (4 units in total) that have been inserted into

10 respective apertures 72 and 82 in the rear of the shelf enclosure 47. The FRUs shown in Figure 4 include two Combined Switch and Service Processors (CSSPs) 71 and two Power Supply Units (PSUs) 81. Various features shown in Figure 4 will be described later.

15 Before proceeding with a more detailed description of each of the FRUs 43, 71 and 81 introduced so far and of the construction of the shelf 41, there follows a brief description of an information processing cartridge 43, a CSSP 71 and a PSU 81 with reference to Figures 3, 4, 5, 6 and 7.

20 Figure 5A provides a perspective view, partly from the rear, of an information processing cartridge 43. Figure 5B provides a perspective view, partly from the front of the same information processing cartridge 43. Figure 5C provides an exploded perspective view of the construction of the information processing cartridge 43. Here is it to be noted that the term "rear" is applied in the context of the position, when

25 installed, of the information processing cartridge, with respect to the shelf 41 (i.e. in this case the "rear" of the information processing cartridge 43 is the innermost part of the information processing cartridge when it is inserted in the shelf 41). Likewise, "front" refers in the present context to the outermost part of the information processing cartridge when it is inserted in the shelf 41.

With reference to Figures 3, 5A and 5B, it will be noted that information processing cartridges are three-dimensional, having a height (h), width (w) and depth (d). If, as in the present example, the information processing cartridges are to be
5 arranged in a one-dimensional array (a row) across the shelf, then efficient packing for the information processing cartridges is achieved where one dimension (here the width, w) is smaller than the other dimensions (here the depth, d, and the height, h). In a particular example, the enclosure of a processing cartridge 43 has a height h, width w and depth d of 115mm, 26mm and 315mm, respectively, although of course,
10 other examples may have different dimensions.

It will be noted that an enclosure 101 of the present example of an information processing cartridge 43 has six, generally rectangular, faces. For convenience only, the face that is visible from the front of the racking when an information processing
15 cartridge 43 is mounted in the shelf 41 is known as the front face 102. The opposite face is known as the rear face 103. In the present example these two faces, as well as top and bottom faces 104 and 105, and the side faces 106 and 107 have the shape of elongate rectangles.

20 Although in the present example the information processing cartridges have six generally rectangular faces, it will be appreciated that other examples could have other configurations. For example, rather than having a generally rectangular shape, the side faces of another example of an information processing cartridge could have the general shape of a triangle (whereby the information processing cartridge may then
25 only have five faces), a pentagon (whereby the information processing cartridge may then have seven faces), and so on. Indeed, one or more or all of the edges could be curved. However, it will be appreciated that the present configuration provides advantages for example, in terms of manufacturing, engineering and packing density within a shelf 41.

In this example, the information processing cartridge enclosure 101 is fabricated from two housing portions. The first housing portion 109 can be fabricated from, for example, a metal (e.g., pressed steel) and can include one side face 107, the rear face 103 and part of the top face 104. The part of the top face formed by the metal portion 109 is given the reference 1041 in Figure 5A. The second housing portion 110 can be fabricated from, for example, a plastics material and can include the other side face 108 and the front faces 102 and bottom face 105 and the remaining part 1042 of the top face 104. In the present example, a plastics material used is PolyCarbonate Acrylonitrile Butadiene Styrene (PCABS), however many other plastics materials such as other ABSs and nylons may be used. The metal portion 109 also includes a part 1021 corresponding to the front face 102 and a part 1051 corresponding to a part of the bottom face 105. The parts are covered by the plastics portion 110 when enclosure 101 is fully assembled. These parts are visible in the exploded view of Figure 5C. The two housing portions 109 and 110 are secured to one another by fixings such as screws 118.

As shown in Figure 5A, a groove 108 runs along the length of the plastic top face portion 1042. This groove 108 is provided for interfacing with a guide member of the shelf chassis (not shown in Figure 5A, but see Figures 8B and 8C). A similar groove 113 is formed in the bottom face 105 of the plastics portion 110 as shown in Figures 5B and 5C. This groove 113 is provided for interfacing with a guide member of the shelf chassis (not shown in Figures 5B and 5C, but see Figure 8D).

The provision of the enclosure 101 means that the information processing cartridge 43 can safely be handled by an operator who is not a skilled technician. Also, through the use of the enclosure 101, the information processing cartridge is a robust unit that protects its inner workings from the outside environment and vice versa. The use of a conductive enclosure, e.g., a metal enclosure, can result in the

information processing cartridge including its own electromagnetic shielding. To this end, the enclosure 101 can also be provided with an internal electromagnetic interference (EMI) shielding member 116 as shown in Figure 5C. The EMI shielding member can be located between the internal components 112 of the information processing cartridge 43 (not described in detail with reference to Figures 5A to 5C, but see Figures 11 and 18 below) and the plastics portion 110. The EMI shielding member can be secured to the plastics portion 110, for example by heat bonding or an adhesive. In other examples, the plastics portion could have a conductive layer deposited thereon, or conductive plastics material could be used. In this example, the EMI shielding member 116 is provided with electromagnetic interference (EMI) fingers 114 to ensure good contact with the shelf chassis and the adjacent components. These fingers 114 extend through EMI finger apertures 1141 in the plastics portion 110.

As shown in Figures 3, 5B and 5C, the information processing cartridge 43 incorporates an injector/ejector handle 111 on the front face 102 for facilitating insertion and latching of the information processing cartridge 43 within an aperture in the shelf 41.

As shown in those Figures, the handle 111 of the injector/ejector lever extends substantially the whole height of the front face of the information processing cartridge 43, thereby increasing the mechanical advantage and facilitating injection and ejection of the information processing cartridge 43. As is further shown in those Figures, the front face 102 of the information processing cartridge 43 has perforations 115, in the present example slits, to allow for airflow into the information processing cartridge 43. The front face part 1021 of the metal portion 109 has perforations corresponding to those in the front face 102 such that airflow into the information processing cartridge 43 is not impeded. It will be noted in Figures 3, 5B and 5C that the handle 111 is narrower in its middle than at its ends. This reduces any masking effect of the

handle 111 on the airflow to the perforations 115 in the front face of the information processing cartridge 43 and facilitates grasping of handle 111. The handle 111 can be bowed out from the front face of the information processing cartridge to further facilitate grasping thereof and to reduce any masking effect with regard to the
5 perforations. In other examples, the handle 111 could have an open frame-like structure to further facilitate airflow.

As shown in Figure 5A, the rear face 103 of the information processing cartridge 43 also has perforations 117 to allow for air to be exhausted from the rear of
10 the information processing cartridge 43. A fan can be located within the enclosure 101 of an information processing cartridge 43. In the present example the fan is a combined radial fan (blower) and heat sink to direct cooling air onto a processor of the information processing cartridge, which fan also encourages air movement within the enclosure 101. LED indicators 119 (see Figure 3 and 5B) can be provided on the front
15 face 102 of an information processing cartridge to indicate whether power is on, whether service intervention is required and whether the information processing cartridge 43 can be removed. A connector 120, for example a 40 way single connector attachment (SCA-2) connector (a small computer systems interface (SCSI) connector), can be provided at the rear of the information processing cartridge 43 for
20 electrical connection of the information processing cartridge 43 within the shelf 41. The connector 120 is advantageously able to withstand repeated removals and insertions of the enclosure 101 from and into a shelf 41. The connector arrangement can include a guide pin arrangement to prevent module misalignment during insertion of the information processing cartridge into the receiving location.

25

Thus an example of features and the construction of an information processing cartridge enclosure has been described. Although particular materials and constructions have been described, it will be appreciated that other examples could be employed. Indeed, it will be appreciated that this particular example relates to only a

possible form for the processing cartridge 43. An example of an alternative construction will now be described with reference to Figure 5D.

Figure 5D provides a perspective view, partly from the rear, of an alternative
5 information processing cartridge 43. Here it is to be noted that the term "rear" is applied in the context of the position, when installed, of the information processing cartridge, with respect to the shelf 41 (i.e. in this case the "rear" of the information processing cartridge 43 is the innermost part of the information processing cartridge when it is inserted in the shelf 41).

10

In this example, the information processing cartridge enclosure 101 is fabricated from pressed steel to form two chassis portions. The first portion 234 includes one side face 107, and part of each of the front and rear faces 102 and 103 and the top and bottom faces 104 and 105. The second portion 235 includes the other
15 side face 108 and the remaining part of each of the front and rear faces 102 and 103 and the top and bottom faces 104 and 105. The two chassis portions 234 and 235 meet at a groove 236 and are secured to one another by fixings (e.g., one or more screws, not shown). Grooves 236 run along the top and bottom faces 104 and 105 of the enclosure 101 and are provided for interfacing with guide rails of the shelf chassis
20 49 (not shown in Figure 5, but see Figure 8C). A cover portion that is secured to the chassis portion forms the other side face 106. It will be appreciated however, that in another example, the chassis portions could be joined at a position other than the groove 108, with the groove 108 being formed entirely in one of the chassis portions. Alternatively, the enclosure 101 may be constructed from a number of sheets of steel,
25 with each sheet forming one of the faces.

Figure 6 provides a perspective view, partly from the front, of a Combined Switch and Service Processor (CSSP) cartridge (or CSSP) 71. Here it is to be noted that the term "front" is applied in the context of the position, when installed, of the

CSSP cartridge 71, with respect to the shelf 41 (i.e. in this case the “front” of the CSSP cartridge 71 is the innermost part of the CSSP cartridge 71 when it is inserted in the shelf 41).

5 With reference to Figures 4 and 6, it will be noted that a CSSP cartridge 71 is three-dimensional, having a height (h), width (w) and depth (d). In a particular example, the enclosure of a CSSP 71 has a height h, width w and depth d of 43mm, 202mm and 278mm, respectively, although of course, other examples may have different dimensions.

10

 An enclosure 121 of present example of a CSSP cartridge 71 has six, generally rectangular, faces. For convenience only, the face that is visible from the rear of the racking when a CSSP cartridge 71 is mounted in the shelf 41 is known as the rear face 122. The opposite face is known as the front face 123. In the present example these
15 two faces, as well as side faces 126 and 127 have the shape of elongate rectangles. The top and bottom faces 124 and 125 are also rectangular, but not elongate in the manner of the front, rear, top and bottom faces. Although in the present example the CSSP cartridges have six generally rectangular faces, as for the information processing cartridges 43 it will be appreciated that other examples could have other
20 configurations.

 In this example, the CSSP enclosure 121 is fabricated from steel sheet to form a chassis portion that includes the bottom face 125, the front and rear faces 122 and 123 and the side faces 126 and 127. A cover portion that is secured to the chassis
25 portion forms the other top face 124. The cover portion is secured to the chassis portion by suitable fixings, for example one or more screws 128. It will be appreciated however, that in another example, other faces, or portions, of the enclosure could form the chassis and the cover portions. The provision of the enclosure 121 means that the CSSP cartridge 71 can safely be handled by an operator

who is not a skilled technician. Also, through the use of the enclosure 121, the switch cartridge is a robust unit that protects its inner workings from the outside environment and vice versa. The use of a conductive enclosure, e.g., a metal enclosure, means that the CSSP cartridge includes its own electromagnetic shielding. To this end the CSSP enclosure 121 is provided with EMI fingers 129 to ensure good contact with the shelf chassis and the adjacent components.

As shown in Figure 4, the CSSP cartridge 71 incorporates two D-shaped handles 132 to facilitate insertion and removal of the CSSP cartridge 71 with respect to an aperture 72 in the rear face of the shelf enclosure. A latch member 131 can be pivotably mounted on a plate that can be secured (e.g., using screws) to the rear face of the shelf enclosure. The latch member 131 is configured to engage one of the handles 132 and to secure the CSSP cartridge 71 in place. In other examples, the CSSP cartridge 71 could be provided with an injector/ejector handle in a manner similar to the information processing cartridge. As shown in Figure 6, the front face 123 of the CSSP cartridge 71 has perforations 133 to allow for airflow into the CSSP cartridge 71. As shown in Figure 4, the rear face 122 of the CSSP cartridge 71 has perforations 135 to allow for air to be exhausted from the rear of the CSSP cartridge 71.

20

At least one fan can be located, for example behind the perforated portion 135 of the rear face, in a CSSP cartridge 71 to channel cooling air through the CSSP cartridge 71 from the front to the rear. In this particular example shown, two fans are provided, one behind each set of perforations 135. LED indicators 137, as shown in Figure 4, can be provided on the rear face 122 of the CSSP enclosure 121 to indicate whether power is on, whether service intervention is required and whether the switch can be removed. Additional link status indicators can be provided integral to 2x4 stacked RJ-45 connectors 139, also shown in Figure 4. As shown in Figure 4, electrical connections 141 can be provided at the front face of the CSSP (i.e. on the

face that in use is inside the shelf enclosure 47). Suitable connections for use in the present example include a connector for power connections, a connector for serial management data connections and a connector for information connections. In the present example, information connections are implemented using an Ethernet
5 information communication protocol, e.g. at 1 Gigabit (Gb). However other protocols could equally be used, for example the Infiniband information communication protocol. The connector arrangement can include a guide pin arrangement to prevent module misalignment during insertion of the CSSP module into the receiving location. For this purpose, guide pin holes 142 can be provided on the front face 123
10 into which guide pins may pass to aid module alignment.

In the present example, up to two CSSPs 71 can be mounted at any one time at the rear of the shelf unit in corresponding apertures 72 in the rear face of the shelf enclosure 47. The number of CSSPs 71 provided in any particular implementation
15 depends upon system configuration and the need, or otherwise, for redundancy.

It will be appreciated that one possible construction of the CSSP cartridge 71 has been described and that as for the information processing cartridge 43, other examples could employ other materials and/or constructions.
20

Figure 7 provides a perspective view, partly from the front, of a power supply unit (PSU) cartridge 81. Here it is to be noted that the term "front" is applied in the context of the position, when installed, of the PSU cartridge 81, with respect to the shelf 41 (i.e. in this case the "front" of the PSU cartridge 81 is the innermost part of
25 the PSU cartridge 81 when it is inserted in the shelf 41).

With reference to Figures 4 and 7, it will be noted that PSU cartridge 81 is three-dimensional, having a height (h), width (w) and depth (d). In this particular example, the order to provide for a dense packing of the FRUs in the rear of the shelf

41, the PSU cartridge 81 has two dimensions (hereinafter described as the width, w, and the depth, d) that are generally similar. In a particular example, the enclosure of a PSU cartridge 81 has a height h, width w and depth d of 83mm, 202mm and 276mm, respectively, although of course, other examples may have different dimensions.

5

An enclosure 145 of present example of a PSU cartridge 81 is of generally oblong shape, but has the "top" "front" edge cut away to form an additional "top" "front" sloping face. The enclosure 145 therefore has five, generally rectangular, faces and two faces of generally rectangular shape with one corner cut away. For
10 convenience only, the face that is visible from the rear of the racking when the PSU cartridge 81 is mounted in the shelf 41 is known as the rear face 146. The opposite face is known as the front face 147. In the present example these two faces and the two side faces 150, 151 are of elongate, generally rectangular shape with one corner cut away, given that the width and depth of the PSU cartridge are similar, whereas the
15 top and bottom faces 148, 149, although still rectangular, are not, in this example, notably elongate. A top front face 148a is present at the top front of the enclosure. Thus the front of the enclosure is sloped at the top edge. As for the information processing cartridges 43, however, it will be appreciated that other examples could have other configurations.

20

In this example, the PSU cartridge enclosure 145 is fabricated from steel sheet to form a housing portion that includes the bottom face 149, the side faces 150 and 151 and the front and rear faces 146 and 147. Cover portions that are secured to the housing portion form the top face 148 and top front face 148a. The cover portions are
25 secured to the chassis portion by suitable fixings, for example one or more screws 152. It will be appreciated however, that in another example, other faces, or portions, of the enclosure could form the chassis and the cover portions. The provision of the enclosure 145 means that the PSU cartridge 81 can safely be handled by an operator who is not a skilled technician. Also, through the use of the enclosure 145, the PSU

cartridge 81 is a robust unit that protects its inner workings from the outside environment and vice versa. The use of a conductive enclosure, e.g., a metal enclosure, means that the PSU cartridge includes its own electromagnetic shielding. To this end the PSU enclosure 145 is provided with EMI fingers 153 to ensure good
5 contact with the shelf chassis and the adjacent components.

As shown in Figure 4, the PSU cartridge 81 incorporates two D-shaped handles 156 to facilitate insertion and removal of the PSU cartridge 81 with respect to an aperture 82 in the rear face of the shelf enclosure. A latch member 155 can be
10 pivotably mounted on a plate that can be secured (e.g., using screws) to the rear face of the shelf enclosure. The latch member 155 is configured to engage one of the handles 156 and to secure the PSU 81 in place. In other examples, the PSU 81 could be provided with an injector/ejector handle in a manner similar to the information processing cartridge. As shown in Figure 7, the front face 147 of the PSU cartridge 81
15 has perforations 157 to allow for airflow into the PSU cartridge 81. As shown in Figure 4, the rear face 146 of the PSU cartridge 81 also has perforations 159 to allow for air to be exhausted from the rear of the PSU cartridge 81.

A pair of fans can be located behind the perforated portions 159 of the rear
20 face of a PSU cartridge 81 to channel cooling air through the PSU cartridge from the front to the rear. LED indicators 161 can be provided on the rear face 146 of the PSU enclosure 81 to indicate whether input power is good, whether output power is good, whether service intervention is required and whether the PSU can be removed. Electrical connectors 163 can be provided at the front face of the PSU (i.e. on the face
25 that in use is inside the shelf enclosure 47) for connection to the shelf. The PSU 81 of the present example may suitably employ an SSI-MPS (Server Systems Interface - Midrange Power Supply) compliant right angle connector at the front face 147 of the PSU 81 to connect to the shelf 41. The power inlet 83 for each PSU 81 can incorporate a cable/connector retention mechanism (not shown) on the rear face 146

of the PSU to prevent accidental or malicious removal of the power input cord from the PSU 81.

In the present example, the shelf unit enclosure provides slots at the rear of the shelf unit for two hot-swappable, AC input PSUs 81. To provide redundant power supply, both PSU slots are populated. More generally, $N+M$ power supplies can be provided, where N is the minimum number of power supply units needed to support the components mounted in the shelf, and M is a selectable number of 0, 1 or more power supplies to provide for redundancy. In the present example, $N = M = 1$.

10

It will be appreciated that one possible construction of the CSSP cartridge 71 has been described and that as for the information processing cartridge 43, other examples could employ other materials and/or constructions.

15 If the full complement of information processing cartridges or switches are not fitted to the enclosure, then blanking panels/modules (e.g., the blanking panels 44 shown in Figure 3) are fitted to all unpopulated positions to maintain the correct airflow and thermal characteristics of the shelf, a safe internal operating temperature for the shelf life expectancy of all shelf components, electromagnetic compliance
20 (EMC) containment and electrostatic discharge (ESD) containment.

It will be noted that each of the FRUs, such as the information processing cartridges 43, is advantageously contained in its own robust enclosure to facilitate EMC containment, ESD containment, handling, storage and transportation. Each
25 FRU can be configured as a 'sealed' unit in the sense that it can be configured not to have field or customer serviceable parts internally. The FRUs can be configured readily to plug into the shelf enclosure and to be hot swappable. The FRUs can be keyed to prevent incorrect positioning and insertion into the shelf enclosure and are arranged positively to be retained in the shelf by a latching/locking mechanism.

The examples of FRUs described above are not provided with removable media. In the present example, internal data storage is provided by 2.5" IDE 9.5mm or 12.7mm profile hard disk drive (HDD) devices mounted internally in each information processing cartridge 43 and in the CSSP cartridge 71. The drives are not considered as FRUs and are not hot-swappable disk drives in the present example, although they could be in other examples. In other examples, the information processing cartridges can be configured without internal hard disk drives.

10 The internal configuration of the shelf 41 and a midplane 171 contained therein is described in the following with reference to Figures 8A, 8B 8C and 8D, and Figures 9A, 9B and 9C.

Figure 8A is a schematic plan view showing the internal configuration of an example of a shelf 41 with the cover 61 removed. Figure 8B is a schematic perspective view from above the rear of the chassis portion 47 of the shelf enclosure with the field replaceable units removed. Figure 8C is a schematic perspective view from below the front of the chassis portion 47 of the shelf enclosure with the field replaceable units and the base 51 removed. Figure 8D is a schematic perspective view from the front and above a part of the base 51 of the shelf 41. Figures 9A, 9B and 9C are, respectively, front, top and rear views of the midplane 171. In this example, the midplane is, in use, mounted vertically within the shelf 41 extending across the width W of the shelf 41 at a position approximately half way between the front and the rear of the shelf 41.

25

The vertically mounted midplane 171 extends, in this example, across the shelf 41 and allows for the electrical interconnection of the FRUs. The various apertures in the front and rear faces 57 and 59 of the shelf 41, in combination with the midplane 171, can be provided with guides (e.g., rails 181) and keying e.g., offset connector

positioning for the insertion of the FRUs into the enclosure and midplane 171. The midplane 171 can be a double-sided, or multi-layer printed circuit board (PCB) assembly that can be mounted vertically in a rigid manner within the enclosure. It can carry connectors 175 on a front surface 172 for making electrical connection with
5 corresponding connectors 120 on the information processing cartridges 43. It can also carry connectors 177 and 179 on rear surface 173 for making electrical connection with corresponding connectors 141 and 163 on the CSSPs 71 and the PSUs 81, respectively. Conductive tracks (not shown) on and through the midplane 171 can be provided to interconnect the various connectors. In addition, the midplane can
10 provide connectors for receiving corresponding connectors connected to first and second indicator boards 183 and 184 that each carry a respective set of LED indicators 69. In the present example, the midplane 171 is not configured as a FRU and is not hot swappable. It is perforated to facilitate airflow through the shelf 41. The midplane 171 can include openings 185, which cooperate with openings in the
15 enclosures of the FRUs 43 and 81, to provide a path for cooling air to pass from the front to the rear of the shelf 41, the cooling air being driven by fans in one or more of the FRUs, for example in the PSUs 81, possibly also in the information processing cartridges 43.

20 A plenum chamber floor member 94 can extend horizontally from the front of the midplane 171 to the front face 57 of the shelf enclosure, or chassis 47. The member 94 provides a floor for a plenum chamber 66, which is supplied with air via the apertures 65 in the front bezel and, in the illustrated example, the slot shaped aperture 68 in the front face 57 of the shelf enclosure 47. Although, for reasons of
25 ease of illustration a slot shaped aperture 68 is shown, a plurality of apertures 68 aligned with the blade receiving locations may be provided. The aperture or apertures 68 can serve both as air vents for a flow of air to the plenum chamber 66, and also as latching locations for latching portions at the top of the injector/ejector levers 111 for

the blades shown in Figures 5B and 5C. The top and sides of the plenum chamber are provided by the top cover 61 and side faces 53 and 54 of the shelf enclosure 47.

A plurality of cartridge guides 97 can be provided at the underside of the
5 plenum chamber floor member 94. In the present example, these guides comprise sprung wire members, e.g., of a resilient metal such as spring steel, that are attached to the top surface of the plenum chamber floor member 94 and extend through a plurality of apertures therethrough to result in a row of guides 97 at the underside of the plenum chamber floor member 94. This arrangement is shown in Figures 8B and 8C.
10 In Figure 8B, the sprung wire members 98 are shown attached to the top surface of the plenum chamber floor member 94. In the present example, the sprung wire members 98 are arranged in pairs, such that two guides 97 are provided by each spring clip 98. In Figure 8C, the guides 97 formed by the protruding portions of the sprung wire members 98 are shown at the underside of the plenum chamber floor member 94.
15 Each guide 97 is advantageously positioned so as to interface with the groove 108 in the plastics material in the upper face 104 of a processing cartridge 43 as shown in Figures 5A-5C to aid correct alignment and to facilitate insertion of the processing cartridge during insertion of the cartridge into the shelf 41. The use of the spring clip as a guide 97 also serves to urge the processing cartridge downwards to provide a
20 secure mounting of the processing cartridge 43, to take account of manufacturing and operational tolerances and to assist in insertion of the processing cartridge where an operator does not align this absolutely correctly.

A further row of cartridge guides 99 can be provided at the upper surface of
25 the base 51 of the shelf 41. In the present example, as shown in Figure 8D, these guides 99 have a rail like form, which can be achieved by punching or stamping through the base 51 of the shelf 41. In this example each guide, or rail, 99 includes a pair of upstands separated by an aperture 100 through the base 51. The size of the aperture 100 can correspond to the width between the upstands. The separation of the

upstands is selected so that the overall width of the resulting rails is slightly less than the width of a groove formed in the lower face of an information processing cartridge 43. Thus, each guide 97 is advantageously arranged so as to interface with the groove 1113 in the plastics material in the lower face 104 of a processing cartridge 43 as shown in Figures 5A-5C to aid correct alignment and to facilitate insertion of the processing cartridge during insertion of the cartridge into the shelf 41.

In the present example, where the guides 97 and 99 are formed from metal, the provision of the grooves 108 and 113 in plastics material at the upper and lower faces, respectively, of each information processing cartridge 43 (see Figures 5A-5B) results in a combination of metal and plastics materials that can provide a low friction interaction, facilitating insertion of the information processing cartridges.

If, for example, the information processing cartridge enclosure is made of a metal, it may be undesirable to provide metal guides to avoid a metal to metal interaction. In such a case, for example, it may be desirable to form the guides from a plastics material having a low coefficient of friction, such as polytetrafluoroethene (PTFE) or polythene. Plastics rails could be attached to the underside of the plenum chamber floor member 94 and/or on the upper surface of the base 51 of the shelf 41. In such an example, grooves on the upper and lower faces of the information processing cartridges 43 could then be formed of metal or plastics and still result in a low friction arrangement.

A CSSP/PSU divider 96 can be provided to the rear of the midplane 171 and can extend horizontally to the rear face 59 of the shelf enclosure 47. The CSSPs 71, when inserted, are supported by the divider 96. To aid the correct insertion of the CSSPs 71, CSSP guide pins 178 are provided on the midplane 171 at positions adjacent connectors 177 on the midplane 171 for connection to the CSSPs 71.

Respective positions 88 and 89 can be formed in the front face 57 and the rear face 59 at which first and second indicator boards 183 and 184 supporting the indicator LEDs 69 can be located. These positions 88, 89 therefore include an aperture through the respective face of the shelf enclosure 47 such that indicator LEDs 5 69 mounted onto a circuit board attached to the inside of the shelf enclosure 47 may be viewed from outside the shelf enclosure.

There now follows are more detailed description of the midplane 171.

10 As mentioned above, the midplane 171 connects all the elements of a shelf together, including, in the present example, up to sixteen information processing cartridges 43, up to two CSSPs 71, two PSUs 81 and the two indicator boards 183 and 184. In the present example, due to its location within the shelf enclosure, the midplane 171 is not configured to be swappable. Accordingly, to maximize the 15 system reliability, the midplane is configured to provide as a high level of reliability as possible. To this end, the midplane is advantageously configured without active devices and to include the minimum number of decoupling capacitors consistent with good design practice (ideally zero).

20 The midplane supports a number of paths for various power and signal lines to interconnect the FRUs.

In the present example, each information processing cartridge 43 has a high speed information signal connection (e.g., a Gigabit (Gb) Ethernet 25 SERIALIZER/DESERIALIZER (SERDES) connection) to each of the CSSPs 71, each connection consisting of two pairs of differential signals. In a conventional manner therefore, the tracking of the paths for these signals is arranged to keep the pairs well balanced and on a single signal layer (i.e. without vias) to support such differential signals at high frequency.

In addition, in the present example, each information processing cartridge 43 has a serial console connection to the CSSP cartridge 71. Each connection consists of two TTL (Transistor-Transistor Logic) level signals that make a transmit and return
5 (TX and RX) pair.

Also, each PSU 81 has a management signal connection (e.g., a serial I2C (Inter-IC Bus) connection) to the CSSP cartridge 71 to control power and monitor environmental parameters. The I2C bus comprises of two signals SCL and SDL
10 (serial clock line and serial data line). In addition, an I2C address programming pin is provided for the PSUs 81.

Each information processing cartridge 43 and PSU 81 can signal to the CSSP cartridge 71 that it is inserted by pulling to ground (GND) a respective Inserted_L
15 signal (i.e., an active low signal). These signals are fed to the CSSP cartridge 71 via the midplane 171.

Each PSU 81 has five 12 Volt output rails. The routing from each PSU 81 is arranged so that a fault in any single FRU cannot completely interrupt the power to
20 any other.

As mentioned above, the midplane 171 is provided with appropriate connector arrangements for receiving the connectors on the FRUs.

25 In the present example, the information processing cartridge 43 connects to the midplane 171 through a 40 pin Single Connector Attachment (SCA-2) connector as defined by the Small Computer Systems Interface (SCSI) standard. Accordingly, the midplane carries corresponding connectors 175.

In the present example, each CSSP cartridge 71 connects to the midplane 171 through a two right-angle 20 pair connector (e.g., 2mm HM-Zd connectors available from Tyco Electronics). The corresponding connectors 177 on the midplane are straight male parts with a power connector. A guide pin arrangement is provided in addition to the connectors to prevent misaligned modules causing bent pins during insertion. The guide pin also provides a leading ground. The CSSP cartridge 71 also connects to the midplane 171 through a right-angled 125 way 5 row 2mm connector. The connector 177 on the midplane 171 includes a straight male part. A guide pin arrangement is provided in addition to the connectors to prevent misaligned modules causing bent pins during insertion.

In the present example, as mentioned above, each PSU 81 connects to the midplane 171 through an SSI-MPS specification connector. The contacts are configured 5P/24S/6P with sequenced signal (S) and power (P) pins. Where the connector on the PSU is a 1450230-1 R/A male header, solder tails connector, the mating connector 179 on the midplane can be a 1450540-2 vertical receptacle, press-fit connector.

In the present implementation, indicator boards 183 and 184 (see Figure 8A) are provided at the front and rear of the system and are configured as FRUs. In this example they hold three system-level indicator LEDs 69 and include a FRU identity (FRU-ID) programmable read-only memory (PROM) each. Three LEDs 69 are present on the indicator board. There can, for example, be a white locator LED that can be switched by the user for locating the system; a green power-on LED to indicate when the system is powered; and an amber service-required LED to indicate a fault or other condition requiring servicing. These LEDs can be driven by the CSSP 71.

In the present example, identification information (FRU ID) for the midplane 171 is held on an I2C electrically erasable programmable read only memory

(EEPROM) in the front indicator board 183. In addition to the I2C signals necessary to access the FRU ID EEPROM, the CSSPs 71 provide a current limited supply to the indicator boards 183 and 184 via the midplane. The indicator boards 183 and 184 are also provided with an I2C address programming pin. Depending on the
5 implementation, FRU ID information can be stored instead, or in addition, on the rear indicator board 184.

As the FRU-ID for the midplane 171 is held on one or both of the indicator boards 183 and 184, the midplane can be a totally passive unit. The FRU-ID PROMs
10 communicate with the CSSPs 71 via an I2C bus. Each device on the bus has a separate I2C address. The lower three I2C address bits of the EEPROMs used are available as pins on the device, to allow programming with resistors. The least significant bit of this address (A0) is passed to the midplane via the corresponding connector. This allows the midplane 171 to program the address of the FRU-ID
15 differently for the front and rear indicator boards 183 and 184, by pulling the address low for the front board and high for the rear indicator board 183. This ensures that both EEPROMS are available on the bus, at different addresses. The FRU-ID for the midplane can be stored on either front or rear EEPROM, but the present example the FRU-ID is stored in the EEPROM on the front indicator board 183. The EEPROM
20 can be 8kByte or larger.

As mentioned above, the midplane 171 includes openings 185 to provide a ventilation path for cooling air passing through the shelf 41. The cooling air passing through the shelf 41 via the midplane 171 can be driven by means of fans provided in
25 each of the information processing cartridges 43 and the power supply modules 81. The openings 185 shown in Figures 8B, 9A, 9B and 9C form schematic representations of openings in the midplane 171. In practice, the openings could have any form (i.e., a series of large openings, or a number of small perforations), arranged on the midplane to align with corresponding openings or ventilation apertures in the

various field replaceable units 43, 71 and 81. In this way, the path of the airflow from the front of the shelf to the back of the shelf can be configured to be as efficient as possible, depending on the detail configuration of the fan units and the ventilation openings or apertures in the information processing, switch, service processor and power supply unit modules 43, 71 and 81. Providing the fan units in the field replaceable units 43, 71 and 81, contributes to the aim of maintaining the chassis 49 and the midplane 171 of the shelf 41 free of active components, thereby minimising cost, and facilitating maintenance. Also, by providing the fan units in each of the field replaceable units, merely inserting and removing field replaceable units automatically adapts the flow of cooling air to the number and type of field replaceable units inserted in the shelf 41.

As described above, in the present example each of the FRUs is designed to be a non-user serviceable unit. Thus each FRU presents the user with a “sealed” unit which may be inserted into and removed from the shelf 41 as desired or required. If a FRU ceases to be operable, then the user has a choice only of returning the FRU to a supplier or service company for repair or of discarding the non-operable unit. As the FRUs are non-user serviceable, there is no requirement for a skilled technician to be employed in inserting or removing the FRUs into or from a shelf 41. Thus each FRU is designed such that a non-skilled person should have difficulty in causing damage to the FRU during handling. Moreover, the configuration and construction of the FRUs (e.g., the injector/ejector levers, the grooves in the enclosures of the information processing units, etc), of the shelf enclosure and the midplane (e.g., the guide rails to guide insertion of the FRUs, the locating pins, etc) contribute to facilitating easy insertion and removal of the FRUs.

Shown in Figure 10 is an example of the flow of cooling air through the shelf 41 and FRUs 43, 71, 81 mounted therein.

In this example, the cooling air passing through the shelf 41 is drawn generally in a front to rear direction through the shelf 41 by cooling fans mounted within the CSSPs 71 and the PSUs 81. Two separate flow paths for cooling air are provided in this example. The first, indicated as flow path $\alpha\alpha$ by dotted lines 77 provides cooling
5 air to the CSSPs 71. The second path, indicated as flow path $\gamma\gamma$ by dotted lines 78 provides cooling air to the information processing cartridges 43 and PSUs 81.

The flow of cooling air along path $\alpha\alpha$ enters the shelf 41 through the aperture(s) 65 in the front face 57 of the shelf enclosure 47 into the plenum chamber
10 66. Further apertures (not shown) could also be provided in the side faces 53 and 55 of the shelf enclosure 47 to allow air to flow into the plenum chamber 66. This air then flows through the plenum chamber 66, and passes over the top edge of the midplane 171 to reach the perforations 133 of the front face of the CSSPs 71. The cooling air then passes through the CSSPs 71, providing cooling to the components
15 thereof before passing out of the CSSPs 71 through the perforations 135 in the rear face of the CSSPs 71 thus being exhausted from the shelf 41. This flow of cooling air along flow path $\alpha\alpha$ is driven by fans 79 mounted within the CSSPs 71. In the present example, a pair of fans 79 is provided within each CSSP 71 and is mounted against the rear face thereof.

20

Air flowing along path $\alpha\alpha$ is impeded from flowing around the processing cartridges 43 by plenum chamber floor member 94 and is impeded from flowing to the PSUs 81 by CSSP/PSU divider 96. This flow path $\alpha\alpha$ therefore ensures that air flowing to the CSSPs 71 is not warmed by passage through the processing cartridges
25 43 and therefore provides maximum efficiency cooling to the CSSPs 71.

The flow of cooling air along path $\gamma\gamma$ enters the shelf 41 through the perforations 115 in the front face of the information processing cartridges 43. The air thus enters the information processing cartridges 43 and provides cooling to the

components thereof. Cooling fans (not shown) within the information processing cartridges 43 direct the cooling air to the processor (CPU) of the information processing cartridge and direct the flow of air in the cartridge thereby increasing cooling efficiency. The air then exits the information processing cartridges 43
5 through the perforations 117 in the rear face thereof. The air then passes through the apertures 185 through the midplane 171 to reach the PSUs 81. This cooling air then passes through the perforations 157 on the front and upper front faces of the PSUs 81 to enter the PSUs and provide cooling to components thereof. It will be appreciated from Figure 10 that the sloping rear of the upper face of the PSUs 81 increases the
10 area over which air can be drawn into the PSUs, thereby reducing the back pressure on the air flowing through the shelf unit and aiding the cooling efficiency. The flow of cooling air along path $\gamma\gamma$ is driven by fans 85 mounted within the PSUs 81. In the present example, a pair of fans 85 is provided within each PSU 81 and are mounted against the rear face thereof.

15

Air reaching the PSUs 81 via path $\gamma\gamma$ will already have passed through the processing cartridges 43. Such air will therefore be already warmed above the ambient temperature outside of the shelf 41 by its passage through the processing cartridges 43. However, as the cooling requirement of the PSUs 81 is typically less
20 than that for the CSSPs 71, this does not cause any difficulty in the operation of the PSUs 81, which are adequately cooled by this flow of pre-warmed air. The pre-warmed air passing through the apertures 185 through the midplane 171 is impeded from flowing into path $\alpha\alpha$ and entering the CSSPs 71 by the SCCP/PSU divider 96.

25 As will be appreciated by the skilled addressee, the arrangement shown in Figure 10 is illustrative only and other arrangements whereby, for example, a mixture of cool air from the plenum chamber 66 and warmed air from the processing cartridges 43 is provided to each rear-mounted FRU can readily be constructed.

With reference to Figure 11, there now follows a description of functional elements of an information processing cartridge 43 as contained within the information processing cartridge enclosure 101.

5 The information processing cartridge 43 includes a microprocessor 192 (a non-limiting example of a microprocessor that can be utilised in the present example is an UltraSPARC™ processor). The microprocessor is mounted on an information processing cartridge motherboard 191.

10 A configurable clock generator 193, configured as a programmable clock synthesizer employing a crystal, can be used to produce CPU clock signals, CLKA and CLKB. The clock frequency can be determined by jumper settings (not shown). A vectored interrupt controller (I-Chip) 194 and a configurable core voltage regulator module (VRM) 195 are provided.

15

 In the present example, memory means for use by the processor 192 when executing instructions can be provided in the form of buffered dynamic random access memory (DRAM), for example configured as dual in line memory modules (DIMMs) 196 with a 72-bit data path with error correction codes (ECC), seated in two sockets
20 on a riser card from the information processing cartridge motherboard 191. The memory capacity can be chosen to suit the processor addressable memory space. For example, in the present example, up to 4 Gigabytes (4GB) of addressable memory can be provided. Serial Presence Detect (SPD) auto-configuration is provided via a Service Management Bus (SMBus) over an I2C bus 197.

25

 In the present example, a PCI bus architecture can be employed with a so-called SouthBridge bus bridge 199 with SuperIO and two Gb Ethernet Media Access Control (MAC) devices. As described above, however, other bus protocols (e.g., Infiniband) can be used. A 32bit PCI bus 198 can be provided from the

microprocessor 192. The SouthBridge 199 is a standard form of bus bridge, in the present example packaged in a 352 pin PBGA (Plastic Ball Grid Array) package, that provides the following functions: an SM Bus interface over the I2C bus 197 for access to the SPD (Serial Presence Detect) feature of the DIMMs that allows
5 initialization of the memory controller; an Xbus interface for access via an Xbus 200 (which is a packet switched multiprocessor bus) to a PROM 201, a real time clock (RTC) 202 and an information processing cartridge service controller (hereinafter termed a Blade Service Controller (BSC)) 203; an IDE (Integrated Drive Electronics) interface that provides an ATA-100 (AT Attachment) IDE connection 204 to an IDE
10 disk drive 205; and a serial console interface on a service bus 206 to the BSC 203 that is used for operating system functions including a console function with this embodiment.

For IO to the midplane 171, two AC-coupled Ethernet interfaces 207 and 208
15 are provided in the present example, which are packaged in a 316 pin PBGA. These Ethernet interfaces can provide a PCI attached Ethernet MAC capable of operation up to Gigabit Ethernet performance. The physical layer can be implemented using SERializer/DESerializers (SERDESS) 209 and 210. An example of a SERDES device is the TLK2201 transceiver manufactured by Texas Instruments, Inc. The SERDES
20 devices use differential PECL TX+/- and RX+/- (Positive Emitter Coupled Logic Transmit and Receive) pairs to communicate to the switch portions of the CSSPs 71 over the midplane 171. The RX+/- pairs can be AC coupled at the information processing cartridge 43, the TX+/- pairs can be AC coupled at each CSSP 71. This facilitates hot-swap of the information processing cartridges 43 and the CSSPs 71.

25

Asynchronous serial connections 211 and 212 for communication between the BSC 203 and the Service Processor parts of the CSSPs 71 can be provided.

Internal data storage can be provided in the present example by a hard disk 205 with a capacity of 30GB or more rated for 24/7 continuous operation. The hard disk 205 is accessed using the primary IDE interface of the SouthBridge 199. The hard disk 205 can hold an operating system, for example a Solaris operating system, and
5 other software and data for performing information processing using the main, or host, processor (CPU) within the information processing cartridge 43.

In the present implementation, the BSC 203 can be implemented as a microcontroller (e.g., a Hitachi H8 microcontroller). The BSC 203 can provide
10 various functions, including for example: dual access (for the information processing cartridges and the CSSPs 71) to PROM 201 and EEPROM 213 for boot information and a FRU-ID for the information processing cartridge; channelling communication between an information processing cartridge 43 and the service processor part of the CSSPs 71; control of power on reset (POR), system reset and externally initiated reset
15 (XIR) to the microprocessor 192; control of the power, service-required and ready-to-remove LEDs 69; upgrading of field-upgradable firmware, via the serial interface; a watchdog function for the operating system; monitoring the speed of a CPU fan 214; and communications with an EEPROM 215 and the operating system via the Xbus 200.

20

In the present example, the BSC 203 can be powered by a 5V service bus (SB) rail as soon as a CSSP 71 and a PSU 81 are fully inserted into the midplane 171, it then turns on other DC/DC converters to provide power to the remainder of the information processing cartridge 43. A BSC reset signal can be derived from a simple
25 conventional power on reset (POR) generator that monitors a 5V supply rail.

In the present example a 1MByte Flash PROM 201 can be provided for storing boot variables for OpenBoot™ PROM (OBP) and Power-On-Self-Test (POST). Further OBP variables can be stored in a second 16kByte (16kB) I2C PROM 215,

accessible via the SouthBridge SM Bus port over the IC Bus 197. The PROM 215 can contain 8kByte for OBP variables and 8kByte of unused space. A 16kByte I2C EEPROM 213 that is accessible via the BSC 203 can contain BSC variables and FRU-ID variables. The EEPROM is nominally divided into 8kByte for FRU-ID and
5 8kByte for the BSC variables. Write protection for the FRU-ID is implemented by BSC firmware. Such write protection may be carried out by, for example, acknowledging instructions to write to the protected area, but not to carry out those write instructions.

10 An environmental monitor sensor 215 can be provided to monitor the CPU and ambient temperatures. This sensor can be accessible via the onboard I2C bus from the BSC 203.

The information processing cartridge 43 can be powered from two, diode
15 commoned, 9V power supply rails 216 and 217. DC/DC converters 218 can be used to provide the voltage levels required by the information processing cartridge 43. The DC/DC converters 218 are supplied by dual 9V inputs 216, 217, individually fused 219, 220 and then diode commoned 221, 222. A 5V DC/DC converter can be turned on as soon as the FRU is fully inserted, with the BSC 203 and required portions of the
20 SouthBridge 199 being powered (the 5VSB rail). A field effect transistor (FET) can be used to gate off the main 5V supply to the rest of the information processing cartridge 43. The DC/DC converter outputs and the main 5V FET can be arranged not to turn on until the BSC 203 turns them on via a signal from the SouthBridge 199. The SouthBridge 199 can be used so that if the BSC 203 is reset (by a watchdog
25 timeout or after a firmware download) the state of the DC/DC converters 218 is not affected. When the remaining outputs from the DC/DC converters 218 are within specification, a PWR_GOOD signal can be asserted low to the BSC 203.

A SouthBridge resume circuit can be operable to run from 3V3, and a simple Zener diode dropper circuit can be used to generate 3V3 from the 5VSB supply.

When the FRU is inserted the inrush current can be limited, for example to
5 <1A, and the rate of rise can be configured not to exceed a predetermined value (e.g., 20A/s) to provide a so-called soft start to facilitate hot-insertion. The intent is to prevent damage to the connectors and to avoid generating noise. A soft start controller 223, which controls a ramping-up of voltage levels, can be enabled when the predetermined signal (Inserted_L signal) is asserted low, this signal is on a short
10 pin in the connector and is connected to ground (GND – not shown) through the midplane 171.

In the present example, a processor impingement fan (processor fan) 214 is configured to run at full speed to cool the information processing cartridge 43 and the
15 fan. The speed of the processor fan and sink can be monitored by the BSC 203, using a tachometer sense pin on the microcontroller. In the event of the fan speed falling below a predetermined speed, or percentage of its nominal speed (e.g., 80%), the BSC 203 can be arranged to issue an alert. The nominal speed of the fan can be recorded as part of the BSC EEPROM contents.

20

The midplane connector 120 for the information processing cartridge 43 is used to establish the connection between the information processing cartridge 43 and the midplane. In the present example it supports up to 84 connections (pins) that will deliver SERDES outputs 224, 225, I2C signals 226, 227, and power 216, 217. Signal
25 connections may be made through a right-angled connector. Power connections may be made through the information processing cartridge right-angled connector. The connector can be configured to facilitate hotswapping of the information processing cartridge, for example by having a low insertion force and/or guide pins to increase the ease of serviceability and prevent module misalignment during insertion.

Interrupts to the processor 192 can be encoded using an encoded interrupt vector mechanism. An I-Chip Emulator (ICE) 228 functions as an interrupt concentrator, receiving all system interrupts and encoding them as an interrupt vector
5 according to an interrupt vector code utilisable by the processor 192. In the present example, where an UltraSPARC™ processor is used, the interrupt vector encoding may be based on a 6-bit interrupt vector code.

With reference to Figure 12, there now follows a description of an example of
10 a combined switch and service processor (CSSP) 71. In the present example, each CSSP 71 provides the functionality of a Switch 73 and of a Shelf Service Processor, or Shelf Service Processor (SSP) 74.

Figure 12 provides an overview of the functional components of the CSSP 71
15 including functional components of the Switch 73 and functional components of the SSP 74. In the present example, most of the components relating to the Switch 73 are mounted on a Switch PCB 231, and the components relating to the SSP 75 are provided on a SSP PCB 232. However, it should be noted that the components located in the lower portion of the switch PCB 321 (i.e., that portion below the SSP
20 PCB 232 as illustrated in Figure 12 logically belong to the SSP 74, rather than to the switch 73. It will be appreciated that such component arrangements are not compulsory for successful operation and that any other component arrangement over any number of component boards can be easily achieved using conventional component arrangement techniques.

25

Firstly, with reference to Figure 12, there follows a description of functional elements of the Switch portions 73 of a CSSP 71 as contained within the CSSP enclosure 121.

The midplane connector 141 on the CSSP 71 establishes the connection between the CSSP 71 and the midplane 171. In the present example, it supports up to 84 connections (pins) that will deliver SERDES outputs 265-268, I2C signals 310, 320, 321 and 322, and power 278, 279. Signal connections may be made through two
5 20-pair right-angled connectors. Power connections may be made through a right-angled connector. The connector can be configured to facilitate hotswapping of the board, for example with a low insertion force. The connector also uses guide pins to increase the ease of serviceability and prevent module misalignment during insertion.

10 A switch microprocessor 240 is provided, in the present example the microprocessor used is a Power PC (MPC8245) packaged in a 352pin Tape Ball Grid Array (TBGA) package. This microprocessor 240 supports between 1MB and 2GB of address space in the present example. It further includes an Embedded Programmable Interrupt Controller (EPIC) that provides 5 hardware interrupts (IRQs) or 16 serial
15 interrupts. There are 4 programmable timers with cascade mode function. DRAM memory for the processor can be provided in the present example by a commodity DIMM 242. The processor 240 can be connected to a 32bit PCI bus 241, which operates at, for example, 33MHz/66MHz.

20 A clock input to the processor 240 can be provided by a clock generator (CLK) 243. The CLK 243 can include a configurable clock generator (not shown) implemented as a programmable clock synthesiser employing a crystal used to produce CPU clock signals. The clock frequency can be determined by jumper settings (not shown). A vectored interrupt controller (I-Chip) (not shown) and a
25 configurable core voltage regulator module (VRM) (not shown) can be provided that operate substantially as described above with reference to the like components of Figure 11.

In the present embodiment two switch ASICs (application specific integrated circuits) 244, 245 are provided (in the present example, BCM5632 Gigabit switch ASICs). Each ASIC can provide twelve GMII Interfaces (1 Gigabit Ethernet) (for uplinks and downlinks) and one 10Gb XGMII interface for chip-to-chip
5 communication (bridging) 246 between the ASICs 244 and 245. Sixteen GMII 1Gb 'downlinks', in the form of serialized Gb Ethernet data, are provided through four quad SERDES 248-251 to allow each information processing cartridge 43 to communicate with the switch 73. Eight GMII 1Gb 'uplinks' are provided for external communication through two quad PHYs 253 and 254 (in the present example
10 BCM5404 ASICs) and RJ45 connectors on the rear panel 122. The ASICs 244 and 245 are configured via a PCI interface (32bit/33MHz) to the PCI bus 241.

A Flash PROM 256 can store a real time operating system, and management and configuration data for the microprocessor. The Flash PROM 256 in the present
15 example can be operable to hold 8MB - 16MB of data, depending on the software required. The flash PROM 256 can be operated via an on-chip XBus 258.

Also connected to communicate with the processor 240 via the XBus 258, a Real Time Clock (RTC) 259 can be provided for real-time functions with a back-up
20 battery.

Also connected to the XBus 258 can be a UART (Universal Asynchronous Receiver Transmitter) 260 which in turn connects to a serial bus 261 for providing an asynchronous console connection from the switch 73 to the SSP 74 which can be
25 accessed by the SSP.

An integrated MAC/PHY (Media Access Control/Physical) switch 271 can provides its own interface to the PCI bus 241. This MAC/PHY switch 271 can connects to a 10/100 Ethernet hub 272. The hub 272 can be operable to provide a

management interface to the SSP 74 and a connection from an external management network to the switch 73 and SSP 74 of a given CSSP 71. The connection from the integrated MAC/PHY device 271 to the SSP 74 can be coupled capacitively. A loopback mode can be provided by the MAC/PHY device 271 for system diagnostics.

5 The hub 272 can connect to an RJ45 connector 273 on the rear panel 122 of the CSSP enclosure 121.

An 8kByte I2C EEPROM 262 can be used to store the FRU-ID and is accessible by the SSP portion 74 of each CSSP 71 via a serial bus 263 and the

10 midplane 171. The upper 2kByte of the EEPROM 262 can be configured to be write protected.

An I2C Redundant Control Register (RCR) 275 can be used to provide an alternate, redundant path for powering-down the CSSP 71 and Shelf Level Indicators

15 69 mounted on the front 57 and rear 59 panels of the shelf 41. The I2C RCR 275 can be accessible by both the SSP 74 of the CSSP 71 containing the RCR and the SSP 74 of a further CSSP 71 connected via the midplane 171 via an I2C bus 276. In the present example, a device suitable for use as the RCR 275 is a Phillips PCF8574 IC.

20 With continued reference to Figure 12, there now follows a description of functional elements of the Shelf Service Processor (SSP) portion 74 of a CSSP 71 as contained within the CSSP enclosure 121 and provided on an SSP PCB 232.

In the present example, communication between the Switch PCB 231 and the

25 SSP PCB 232 is facilitated by an interboard connector pair 298 and 299. It supports connections (pins) for I2C signals, 10/100 MAC/PHY output, and power. As described above, the switch PCB 231 carries the components associated with the switch, and it also carries the power, FRU-ID and environmental monitoring components along with the connectors for connections to the midplane 171 and

external connectors. Thus, in the present example, all SSP components requiring a connection to the midplane 171 or an external connection have signal paths routed through the connector pair 298, 299 and via the switch PCB 231 to the relevant midplane or external connectors.

5

In the present example, the SSP 74 includes a microprocessor 301 (e.g., a Power PC (MPC8245) processor) mounted on the SSP printed circuit board (PCB) 232. The processor 301 can be connected to a PCI bus 302, the present instance a 32 bit bus that operates, for example, at 33MHz/66MHz.

10

A clock input to the processor 301 can be provided by a clock generator (CLK) 303. The CLK 303 can comprise a configurable clock generator (not shown) implemented as a programmable clock synthesiser employing a crystal used to produce CPU clock signals. The clock frequency can be determined by jumper settings (not shown). A vectored interrupt controller (I-Chip) (not shown) and a configurable core voltage regulator module (VRM) (not shown) can be provided that operate substantially as described above with reference to the like components of Figure 11.

20 The processor 301 can be provided with a DRAM memory 305. The memory capacity can be chosen to suit the processor addressable memory space. In the present example, 8 MB of DRAM memory is provided.

An integrated MAC/PHY switch 306 can provide its own interface to the PCI
25 bus 302. The MAC/PHY switch 271 can be connected to 10/100 Ethernet hub 272 via the interboard connectors 298, 299. A loopback mode can be provided by the MAC/PHY switch 306 for system diagnostics.

Octal UARTs 308 and 309 can be connected between the PCI bus 302 and the interboard connector pair 298, 299. The signal path can be continued from the interboard connector pair 298, 299 to serial connections 310 on the midplane connector 141 on switch PCB 231. The Octal UARTS 308, 309 can facilitate serial
5 communications between the SSP 74 and each of the processing cartridges 43.

Also connected to the PCI Bus 302 can be a dual UART (DUART) 312 that in turn can connect via the interboard connectors 298, 299 to serial bus 261 for providing an asynchronous console connection from the SSP 74 to the switch 73. The DUART
10 312 can also have an I2C connection to an external connector on the rear face 122 of the CSSP enclosure 121. The external connector can provide a common operating system/boot console and command port 311.

Connected to the processor 301 via an XBus 314 can be a Flash PROM 315.
15 The Flash PROM 315 can store a real time operating system, and management and configuration data for the microprocessor 301. The Flash PROM 315 can be operable in the present example to hold up to 2MB of data, depending on the software required.

Also connected to the processor 301 via the XBus 214 can be a real time clock
20 (RTC) 316 for real-time functions with a backup battery. The RTC 316 can also provide 8kByte of non-volatile random access memory (NVRAM), in the present instance implemented as an EEPROM. This can be used to contain information such as the FRU-ID, a serial number and other FRU information.

25 To facilitate I2C communications between the SSP 74 and the other CSSP 71, the midplane 171 and the PSUs 81, a multiplexer 318 can be provided. The multiplexer 318 can have a single I2C connection to the processor 301 and connections, via the interboard connector pair 298, 299 and the midplane connector 141 to both PSUs 81, the midplane 171 and the other CSSP 71.

The processor 301 can also comprise an embedded DUART to provide a redundant serial link to the SSP 74 of the other CSSP 71. Although it would be possible to implement this link using an external DUART, the advantage of using an
5 embedded DUART is that the connection to the other CSSP is reliable and therefore likely to be functional. Where the embedded DUART link does not use the I2C Multiplexer for communications to the other CSSP, a common mode of failure for both the SSP – SSP I2C links can be avoided, it being assumed that the processor 301 is likely to be functional even if both embedded DUART channels are non-functional.

10

The CSSP 71 can be powered from two, diode commoned, 9V power supply rails 278 and 279. DC/DC converters 281 can be used to provide the voltage levels required by the CSSP 71. The DC/DC converters 281 can be supplied by dual 9V inputs 278, 279, individually fused 285, 286 and then diode commoned 287, 288. A
15 soft start controller 283 can be provided to facilitate hot-insertion. A 5V DC/DC converter (I2C power regulator) 282 can be turned on as soon as the CSSP 71 is fully inserted. A 3.3V DC/DC converter can be turned on when instructed, for example through SSP service software, by asserting low an appropriate signal (ON_L – not shown). The 3.3V converter can be arranged to turn on a converted for 2.5V, 1.2V,
20 and a processor core voltage rail (Vcore) when the voltages are within an appropriate range.

When the CSSP 71 is inserted the inrush current can be limited, for example to <1A, and the rate of rise can be configured not to exceed a predetermined value (e.g.,
25 20A/s) to provide a so-called soft start to facilitate hot-insertion. The intent is to prevent damage to the connectors and to avoid generating noise. A soft start controller 283, which controls a ramping-up of voltage levels, can be enabled when the predetermined signal (Inserted_L signal) is asserted low, this signal is on a short pin in the connector and is connected to ground (GND – not shown) through the

midplane 171 until one of the supplies is removed. These circuits can be configured to withstand an overvoltage at their inputs whilst the input they are feeding is not powered, without any leakage to the unpowered circuit. A sense circuit can detect if the voltage has dropped below a threshold, for example 2.0V, as a result of a blown
5 fuse, a power rail going down, etc. The DC/DC converters 281 can be protected against short circuit of their outputs so that no damage occurs.

The I2C regulator 282 can be powered as soon as the CSSP 71 is fully inserted into the midplane 171. This can be facilitated through short pins connected to the soft
10 start controller 283, which controls a ramping-up of voltage levels. The other DC/DC regulators can be turned on, for example by SSP software.

A pair of fans 290, 291 can provide cooling to the CSSP 71. The fans 290, 291 can be configured to run at full speed to prevent overtemperature conditions by
15 minimizing the temperature of the internal components and the fan. The speed of the fans 290, 291 can be monitored by the SSP 74 through an environmental monitor 295 on the switch board 231. The environmental monitor 295 can be alerted in the event of the fan speed falling below a predetermined value (e.g., 80% of its nominal speed). The fan can provide tachometer outputs to facilitate the measurement of fan speed.

20

LED indicators 137 can be provided, for example with a green power LED, an amber LED for indicating that service is required and a blue LED for indicating that the switch is ready to be removed. LED indicators integrated on 2x4 stacked RJ45 connectors on the rear face of the CSSP 71 can be arranged, for example, to show
25 green continually when the link is present and flash green when the link is active.

The environmental monitor ENV MON 295 can be provided to maintain operational integrity of the CSSP 71. The ENV MON 295 can include limit values in limit registers and can monitor, for example, temperature within the CSSP enclosure

121, the CSSP power rails, including the 12V, 3V3, Switch Processor Core Voltage, CSSP Processor Core Voltage and the two 9V power feed rails 278, 279 from the midplane 171. The outputs of the DC/DC converters 281 can be fed in to A/D inputs of the ENV MON 295 for Watchdog comparisons to be made to the voltage limits set
5 in the limit registers. As noted above, the ENV MON 295 can also monitor the operating speeds of the fans 290 and 291. The ENV MON 295 can communicate with the SSP 74 of both CSSPs via an I2C bus 296.

For IO to the midplane 171 shown in Figures 8A-10, the midplane connector
10 141 can include sixteen 1Gb Ethernet connections 265-268 from four quad SERDES 248-251 and the I2C bus lines 596.

The SSP 74 can access the I2C devices (FRU-ID EEPROM, 8-bit I/O expansion chip, and the system hardware monitor) through the midplane 171.
15

For external IO, rear panel Gb Ethernet connections can be provided from the two quad PHYs 253, 254 to 2x4 stacked RJ45 connectors 139 (to give 8 uplinks). Each port can be an independent 10/100/1000 BASE-T (auto negotiating) port. The PHY devices 253, 254 can operate in GMII mode to receive signals from the 8-
20 Gigabit interfaces on the ASICs 244, 245.

The Power Supply Units (PSUs) 81 can be configured such that when two or more PSUs 81 are connected in parallel in the shelf 41, failure of any one of the paralleled units shall not affect system operation. Moreover, one of the PSUs can be
25 installed or removed from a "live" system with or without input power applied. The outputs can have overcurrent protection.

The PSU can have an I2C interface to provide power supply status via the midplane 171. The PSU can have an internal temperature sensor that reports via the

I2C interface. The PSU fan speed can also be monitored and errors are reported via the I2C interface. Overvoltage and overcurrent sensors can also report via the I2C interface.

5 There now follows a description of aspects of an example of a power supply 81 with particular reference to Figure 13.

When a power supply (e.g. mains power, or UPS type protected power) is connected to the cable connector 83, transformer, regulator and rectifier circuitry 400
10 can operate to generate a DC output (in the present example, 12V DC) from the input (in the present example 230/240V 50Hz AC or 110V 60Hz AC).

In order to ensure good cooling reliability within the shelf 41, each PSU 81 can have a pair of cooling fans 402, 403 located at the rear of the PSU enclosure as
15 described above with reference to Figure 7. As a failsafe measure against a failure of one PSU 81, the fans of each PSU 81 can be powered by both PSUs 81. Thus, in the event that one PSU 81 is non-operative to produce the DC supply for the shelf 41, as long as the other PSU 81 is operative, not only will all components of the shelf 41 be powered as normal, the fans of both PSUs 81 can continue to run.

20

As shown in Figure 13, this dual powering of cooling fans 402, 403 can be effected by providing a power supply line 404 from the transformer, regulator and rectifier circuitry 400 to power both fans 402, 403. Also, first and second separate power lines 410, 412 from the other PSU 81 can provide duplicate power supply to
25 the first and second fans 402, 403 respectively. The fan 402 can thus be powered by a diode commoned supply from line 404 and a diode commoned supply from line 410. Diode protection can be provided by diodes 405 and 411 respectively. The speed of the fan 402 can be controlled by a speed controller 408. Similarly the fan 403 can be powered by a diode commoned supply from line 404 and a diode commoned supply

from line 412. Diode protection can be provided by diodes 406 and 414 respectively. The speed of the fan 403 can be controlled by a speed controller 409.

The two speed controllers 408, 409 can in turn be controlled by a data input
5 from each CSSP 71 received via an I2C bus connection (not shown in Figure 13). Power supply lines carrying DC power for the other FRUs of the shelf 41 are shown in Figure 13 as power line 416. All power connections to and from the PSU 81 can connect to the midplane 171 when the PSU is inserted in the shelf 41 via the midplane connector 163. In the present example the PSU 81 connects to the shelf through a
10 5P/24S/6P configuration SSI-MPS compliant right angle connector 163 at the front face 147 of the PSU 81. Connectors for the I2C interface can also be provided.

The input power line 410 and 412 for each fan 402 and 403 can be provided with a softstart module 4131 and 4132 respectively, to allow for hot insertion of the
15 PSU 81 into the shelf 41. The softstart modules 4131 and 4132 can be controlled, for example, by pulling a signal to ground (e.g., a "mated" input line 4151 and 4152).

Where the two input power lines 410 and 412 are separate lines having separate softstart provision, there is no common failure mode for the backup method
20 of powering the fans 402, 403. Thus even if the PSU 81 ceases to be operable to generate the DC supply, and a component (power line or softstart module for example) fails in the supply path from the second PSU 81 to the fans 402, 403, at least one of the fans 402, 403 can remain operational as the shelf 41 still receives the cooling effect of three PSU fans.

25

In the present example, the power supply has four rear panel LED indicators 137. A blue "Ready to Remove" LED can be driven by the I2C interface and indicate that the power supply may be removed from the system. An amber "Service Required" LED can be driven by the I2C interface and indicate that the power supply

is in a fault condition: any output out of range, over-temperature or shutdown. A green "DC Output-OK" indicator can be driven by internal power supply circuits and show that the main 12 volt supply is functioning. The LEDs can remain lighted when individual outputs are in the current limited mode of operation. A green "AC Input-
5 OK" indicator can be driven by internal power supply circuits and show that AC input power is within normal operating range.

With reference to Figure 14, there will now be described an example of data connectivity between the FRUs and midplane of the shelf 41. Power transmission
10 paths are not illustrated in Figure 14. However, it will be appreciated that to facilitate the maximum component redundancy of the shelf 41, each PSU 81 can independently provide power to each FRU.

In the present example each of the processing cartridges (blades) 43 connects
15 to the midplane 171 via a pair of information signal connections (e.g. Gb Ethernet links) 224, 225 and a pair of serial management signal connections 226, 227. Connections within the midplane 171 can ensure that each Ethernet link 224 is directed to a connection 265-268 from the midplane 171 to a first switch 73, and that each Ethernet link 225 is directed to a connection 265-268 from the midplane 171 to a
20 second switch 73. Thus one Ethernet link can be established between each processing cartridge 43 and the switch 73 of each CSSP 71. Further connections within the midplane 171 can ensure that each serial connection 226 is directed to a connection 310 from the midplane 171 to the first SSP 74 and that each serial connection 227 is directed to the second SSP 74. Thus one serial link can be established between each
25 processing cartridge 43 and the SSP 74 of each CSSP 71. As mentioned earlier, information signal connections other than Gb Ethernet connections (e.g. Infiniband connections) could be employed in other examples.

A plurality of serial connections can connect each SSP 74 to the other. Serial lines 320, 321 can connect each SSP 74 to the midplane 171 and connections within the midplane can connect the two sets of lines together. To provide a control interface from the SSPs 74 to the PSUs 81, serial lines 322 can connect each SSP 74 to the midplane 171 and connections within the midplane 171 can connect to serial lines 324 from the midplane 171 to each PSU 81.

An example of data and control connectivity of the shelf 41 to and from computer systems external to the shelf 41 when the shelf 41 is arranged for use within a multiprocessor server system such as that described above with reference to Figure 1 will be described with reference to Figure 15.

As summarised above with reference to Figure 14, in the present example each processing cartridge, or blade, 43 is connected to the switch 73 of each CSSP 71 by an information signal connection (e.g. a 1Gb Ethernet link) formed by a combination of links 224, 225 from the processing cartridge 43 to the midplane 171, connections within the midplane 171 and links 265-268 from the midplane 171 to each switch 73.

Further, in this example a set of serial management signal connections comprising links 320, 321 and connections within the midplane 171 connect the SSP 74 of each CSSP 71 to the SSP 74 of the other CSSP 71.

To provide external data connectivity between the shelf 41 and an external core data network 330, in association with which all information processing performed by the processing cartridges 43 of the shelf 41 is undertaken, connections 331 can be formed between the core data network 330 and the eight 1Gb Ethernet ports 139 provided on the rear panel 122 of the CSSP enclosure 121.

In the present example, the connections by means of which control and configuration of the shelf 41 are performed are entirely separate to the connections to the core data network 330. Therefore, a first external switch 335 can connect to a management (I2C) port 273 of the first CSSP 71 and a second external switch 336 can
5 connect to a management (I2C) port 273 of the second CSSP 72. As described above with reference to Figure 12, the management port 273 can provide a management network interface to both the switch 73 and SSP 74 of each CSSP 71. The external switches 335, 336 can each be connected to each of a pair of System Management Server (SMSs) 338, 339. The SMS is not essential to the operation of the shelf 41,
10 but use thereof aids optimal operation of the shelf 41. In a typical multiprocessor server system a plurality of shelves 41 may be connected together via the core data network 330 under the control of a single management network utilising one set of SMSs 338, 339. A set of SMSs 338, 339 may comprise a single SMS (as well as a plurality thereof). However use of at least two SMSs enables redundancy of
15 components, therefore increasing overall system reliability.

A serial interface control 343 operable under telnet protocol control is also connected to the shelf 41 in the present example. This can provide a common operating system/boot console connection to the SSP 74 of both CSSPs 71 via the
20 RJ45 connector 311 on the rear panel 122 of each CSSP enclosure 121.

It will be appreciated from the above that a flexible and scalable modular computer architecture has been described. In the described example up to 16 information processing cartridges, or blades 43, can be configured as sealed FRUs on
25 a single shelf 41, the number of blades being chosen according to customer requirements. Each blade has its own processor and random access memory. If, for example, there is a maximum of 2Gbytes of memory per information processing cartridge, and one processor per blade, 16 processors (16P) with 5.33 processors per unit height (1U) and a total of 32GB of memory per shelf can be provided.

In the present example, the shelf 41 incorporates redundant combined switch and shelf service processor modules (CSSPs) 71 and redundant power supply units (PSUs) 81 separate from the blades 43. As the power supplies are carried by the shelf, the information processing cartridges can be kept compact and inexpensive. Also, as a result, they can be powered by DC power only, via the midplane 171.

Also, as mentioned earlier, the FRUs (e.g., the information processing cartridges, or blades, 43, the CSSPs 71 and the PSUs 81) can all be configured as sealed units that do not contain any internal FRUs themselves and do not contain user serviceable items. The enclosures of the FRUs can be arranged to enclose all of the functional components of the FRU with only electrical connectors being externally accessible and with indicator LEDs being externally visible as well.

These factors can all contribute to keeping the cost of the FRUs low, as well as that of the overall system. The modular approach with the use of sealed modular field replaceable units for providing system functionality and with non-field replaceable units designed with a minimum possible number of active components enhances reliability. Moreover, easy and rapid maintenance is facilitated in the event of a failure of a FRU by simple replacement of that FRU, further reducing the cost of ownership.

Thus, it will be appreciated from the above description that the provision of a rack mountable shelf, that includes power supplies, a shelf service processor and switches in modular units, for carrying a number of processing cartridges, wherein the number of processing cartridges can be chosen according to customer requirements, provides a flexible and scalable computer configuration. The balancing of the load between the processors of the processing cartridges can be effected by software using conventional principles.

A configuration as described provides an easily scalable processor architecture, whereby the processing power provided by a complete system based on the information processing cartridge/information processing cartridge carrier architecture can be scalable from moderate to very high capacity through the simple addition of further information processing cartridges.

As illustrated in Figure 16, an example of the external connections from a shelf 41 can be in the form of two active information signal connections (e.g., Ethernet connections) 350 and 351, two active power connections 353 and an active/standby pair of management connections 354. With regard to the management connections, each connection comprises a serial connection and a network (e.g. Ethernet or Infiniband) connection. It is possible to connect to either the active or the standby connection, as the incoming signal will be internally routed to whichever management controller (CSSP) is the current master. It will be appreciated, therefore, that the connections to a shelf can be kept to a minimum. It will further be appreciated from the configuration shown in Figure 16 that the system is scalable beyond a single shelf unit 41.

Figure 17 illustrates how a plurality of shelves can be configured within one (or more) racks to provide even higher processing power. Such a constellation of shelves to provide a large grouping of servers is sometimes termed a "web farm" or "server farm" 360. As shown in Figure 17, the web farm comprises a plurality of shelves 41 that each carry a plurality of blades 43. Also provided are a plurality of Network Attached Storage devices (NAS) 373 for providing storage for critical data, e.g. email data storage, for the web farm. The NASs 373 are not required if there is no critical data to be stored, e.g. if the web farm is operating solely to provide web caching services.

Management control of the web farm 360 can be provided through a pair of System Management Servers (SMSs) 362. Each SMS 362 can be connected to a management network via a link 366 and to a management console 365. The SMSs 362 can communicate with the individual shelves 41 via a pair of management
5 switches 364. Each shelf 41 and NAS 373 can be connected to each management switch 364 via a connection 367. Thus dual redundant management connections can be provided to each shelf 41 and NAS 373.

Flow of data to and from the web farm 360 can be provided through a pair of
10 data switches 369. Each data switch 369 can be connected to a consumer network via a link 370. It is to be understood that the consumer network can be a larger data network to which the web farm 360 is connected. This network can be an office or corporation intranet, a local area network (LAN), a wide area network (WAN), the Internet or any other network. Connections between the data switches and the shelves
15 41 can be facilitated by connections 371. It is to be noted that as each shelf has its own switching capability, there is no need for each shelf 41 to be directly connected to the data switches 369. Connections can also be provided to connect the NAS units 373 to the shelves 41. The topology used for interconnection of the data switches 369, shelves 41 and NASs 373 can be any topology providing at least one connection of
20 any length between every possible pair of units. Complex topologies arranged to minimise the maximum connection length between any two given units in the web farm can be used.

The web farm 360 comprising a plurality of shelves 41 with or without a
25 plurality of NASs 373 can suitably be used as any or all of the entry edge server group 9, web edge server group 15 and application servers 19 described above with reference to Figure 1.

As an alternative to providing critical data storage within a NAS 373, such storage can be provided within one or more NAS cartridges fitted into one or more of the shelves 41 in place of processing cartridges 43. Another alternative is to provide a server shelf with local storage such as a RAID array (Redundant Array of Inexpensive
5 Disks) in place of the NAS 373.

Thus there has now been described an example of a fully configurable computing system based on a plurality of self contained field replaceable units (FRUs) and scalable from a single processing cartridge with power supply and switching
10 capability to a multiply redundant multiprocessor server system with full system management capability extending over a number of co-operably connected server shelves. It will of course be readily apparent to the skilled reader that many of the specific features specified in the above description are in no way limiting and a variety of alternatives may be produced using only ordinary skill and common general
15 knowledge. Non-limiting examples of example modifications which may be made to the above described system are discussed hereafter.

There is no limit placed on the processing cartridges as to what software they should run. Each module within a shelf or farm may run under the same operating
20 system, or a plurality of different operating systems may be used. Examples of possible operating systems include Sun Microsystems' Solaris® OS or another UNIX™-Type OS such as Linux™, MINIX™, or Irix™, or UNIX™ or a Microsoft OS such as Windows NT™, Windows 2000™, Windows ME/98/95™, Windows XP™.

25

It is also not necessary that each processing cartridge within a shelf or farm be configured to run the same program software. For example, individual processing cartridges may be configured to execute, for example, fileserver software, mailserver

software, webhosting software, database software, firewall software, or verification software.

Although it has been described above with reference to Figures 4, 8, 12, 13
5 and 14, that functionality of a switch and of a shelf service processor should be provided within a single combined switch and service processor unit, this is not essential and separate switch and shelf service processor field replaceable units may be used.

10 Although it has been described above that a pair of PSUs and a pair of CSSPs may be provided so as to enable dual-redundancy, further PSUs and CSSPs may be provided so as to increase FRU redundancy further, thus providing statistically higher reliability.

15 In the power supply circuitry in each of the blades and CSSPs, two voltage sense circuits may be provided after the fuses and before the diodes, to prevent a latent fault caused by a failed fuse going undetected until one of the PSUs is removed or taken offline. Such circuits may be configured to withstand an overvoltage at their inputs whilst the input they are feeding is not powered, without any leakage to the
20 unpowered circuit.

Although it has been described above with particular reference to Figure 11 that the processing module may be based on an UltraSPARC™ processor, this is not limiting and any other processor having sufficient processing capacity to undertake the
25 tasks required of a particular processing cartridge may be used. Alternative processors include, but are not limited to, Intel x86 series and compatible processors, AMD x86 compatible processors, Alpha processors and PowerPC processors. The particular example of an x86 compatible processor is described in more detail with reference to Figure 18. In Figure 18, the parts corresponding to those of the UltraSPARC™ based

system of Figure 11 have the same reference numerals and will not be described again here. In a system based on an x86 compatible processor, the processor 378 itself communicates with the other components, including the memory 196 and PCI bus 198 via a Northbridge 379. The Northbridge 379 also includes an interrupt controller, so
5 no separate interrupt concentrator is required. The other components of the processing cartridge could be substantially the same as for the UltraSPARC™ based system described above.

Although it has been described above that each information processing
10 cartridge comprises a single microprocessor, this is not a limiting case as each or any of the information processing cartridges may have more than one microprocessor arranged to share common storage resources to operate synchronously (in lockstep) or asynchronously. Also, it is not necessary that all information processing cartridges inserted into a shelf at a given time are identical, rather a variety of different blade
15 architectures may be used simultaneously.

The provision of the functions of both Switch and Shelf Service Processor within a single FRU in the present example provides a facility within a single shelf 41 for dual redundancy in both functions in fewer different FRUs. As will be
20 appreciated, there is no restriction that these two functions must be provided within a single FRU and division of the two functions into separate FRUs would present no difficulty to the skilled addressee.

The backup power provision for the cooling fans of the PSUs is described
25 above with reference to Figure 13. Although it is described that the backup power supply to each fan should be independent of the other, if the risk of common mode failure of backup power supply is judged to be low, or is of low importance, then the backup supply to each fan may be provided in common with the backup supply to all other fans.

Also, an information processing module forming a field replaceable server blade can include a processor and memory can be configured by means of software, firmware or hardware to provide a special purpose function. By way of examples
5 only, an information processing module can be configured to perform the function of one or more of a firewall, or a load balancer, encryption and/or decryption processing, an interface to a secure network, e.g. a virtual private network (VPN), a specialized switch with wide area network (WAN) connectability.

10 Also, a storage blade may be provided. The storage blade can be configured to be mountable in a server blade receiving location in a blade server carrier. The storage blade can comprise storage blade connectors configured for interconnecting with carrier connectors on the server blade carrier, whereby the storage blade is interchangeable with a server blade. A carrier, or shelf, for a server system can be
15 arranged with a plurality of blade receiving locations for receiving blades, wherein the blades can be storage blades or information processing blades. The server system can be self configuring on receipt of the blades according to the type of blade received in each said location. To achieve the blade service controller in each blade can be operable to communicate with a shelf service processor to perform said configuring.

20

One aspect of ensuring maximum possible computer availability relates to service and management of a computer system. In particular, service and management functions are in general required for monitoring system status and performing control actions in response to monitored parameters. There now follows
25 an example of how service functions may be organised to be implemented in a hierarchical manner through different components of a modular computer system.

As described above with reference to Figures 11 and 18, each information processing cartridge 43 has a blade service controller 203. The blade service

controller 203 is a local controller of service functions for each information processing cartridge 43. Each blade service controller 203 is arranged to be able to communicate with the SSP 74 (see Figures 12 and 15) of each CSSP 71 via the connections 226, 227 and 310. The SSP 74 provides shelf-level service functions for the population of each shelf 41, which includes the PSUs 81 and the switches 73 of the CSSPs 71 as well as the processing cartridges 43. Each SSP 74 is arranged to be operable to communicate with an external system management server (SMS) 338, 339 via I/O connector 273 as illustrated in Figure 15. The system management server provides system level service functions for one or more shelves 41. The system management server may also provide service functions to computer system modules other than shelves 41 within a computer system, for example to individual server units and/or to reliable storage units such as Network Addressable Storage (NAS) units.

In the present example, the blade service controller 203 can provide the following service functions. It typically performs environmental monitoring of the information processing cartridge 43 within which it is situated, this includes monitoring of CPU and ambient temperatures by means of the environmental monitor 215, monitoring of the power supply rails from the DC to DC converter 218 and monitoring of the speed of the processor cooling fan 214. The blade service controller 203 may also perform a watchdog function for monitoring the operating system running on the processor of the processing cartridge 43. The blade service controller 203 can also provide an out-of-band interface between the processing cartridge 43 and the CSSP 71. The out-of-band interface is a physically discrete interface which is independent of the operating system and applications running on the processor, this interface may be used for providing processor cartridge parameters to the shelf SSPs 74 at boot. In addition, the blade service controller 203 may also provide local functions for control of the LED service indicators 119 and control of power on reset (POR) and externally initiated reset (XIR) to the processor. In the present example, there is no requirement for the blade service controller 203 to have any level of

intelligence, it being sufficient that it is capable of performing simple watchdog and threshold type monitoring operations. However, the blade service controller 203 may be implemented having intelligence, which may be of particular benefit in the event that it is tasked with more complex functions than those set out above.

5

The SSP 73 (of which two are typically provided within a shelf 41 to provide redundancy) provides the following functions in the present example. It typically performs environmental monitoring for the shelf 41, monitoring the PSUs 81, the CSSPs 71 and the PSU inputs. The SSP 73 can also aggregate data from the blade
10 service controller 203 of each blade 43 present in the shelf. The SSP 73 may also be configured to make decisions regarding aspects such as power control, control of the shelf service indicator LEDs 69 and to provide instructions to the blade service controllers 203. In the present example, the SSP 74 is not responsible for deployment of software to individual processing cartridges, however this function may be
15 provided by the SSP 74 in alternative arrangements.

In order to provide a facility for manual monitoring and control of the SSP 74, it may be configured to have a user interface via which a user can be provided access to the functions and data of the SSP 74. The user interface access may be provided
20 through one of the external serial interface connection 311 and the external network connection 273. The SSP 74 may be operable to operate as an isolated unit with no higher level service control provided. In general the SSP 74 requires little or no user intervention to manage the operation of a single shelf 41.

25 Fault management may be controlled by the SSP 74 such that a failed processing cartridge 43 (the failure of which may be detected, for example, from data provided by a blade service controller 203 indicating a failure of the processing cartridge's operating system or application or from errors detected in the switch 73) may be remotely restarted by the SSP 73, in the present example either by initiating an

externally initiated reset (XIR) to the processor of the processing cartridge, or by causing an interruption in the power supply to the processing cartridge sufficient to cause a power on reset (POR) to the processor of the processing cartridge. In the present example, a failed processing cartridge will be the subject of a predetermined
5 number of restart attempts before that processing cartridge is considered permanently failed, at which time the processing cartridge is removed from service in that it is ignored by the SSP 74 and the switch 73 is instructed to pass traffic to it and to ignore any traffic from it and the failed nature is reported to any higher level service control to which the SSP 74 is connected.

10

In order to perform the above described functions, the SSP 74 of the present example is configured to have a degree of intelligence such that decision making steps and provision of a user interface are supported.

15 The system management server (SMS), of which two (338 and 339) are typically provided as a clustered pair for redundancy (as illustrated in Figure 15), is in the present example configured to provide the following functions. The primary use of the server management system 338, 339 of the present example is to provide aggregate management and monitoring for a number of shelves 41 (as illustrated in
20 Figure 17). By means of such aggregated monitoring and management, a single (clustered pair of) system management server(s) can oversee the running of a large number of computer system shelves 41.

The system management server (SMS) 338, 339 may also be operable to
25 perform hardware fault failover in response to hardware failures detected by the system management server or reported to it by an SSP 74 of a particular shelf 41. In addition, the system management server 338, 339 of the present example may be operable to perform software deployment to individual processing cartridges. In

alternative arrangements, the system management server 338, 339 may be operable to allow platform specific abstraction.

By platform abstraction is to be understood that the SMS 338, 339 may be
5 operable to allow platform specific abstraction such that the user may configure computing services without knowledge of the underlying hardware providing these services. For example a user may require a web server which can cope with, say, 1000 connections per an hour. The SMS 338, 339 may be pre-configured with enough platform specific data to determine how many blade servers are required to
10 fulfil this requirement, based on their processor, memory size, etc and provision the web service on that number of blade servers without the user needing to know anything about the server hardware details himself.

The operation of the system management server 338, 339 may be controlled
15 via a management network 366, such that control may be automated (via a higher level of system management) or manual (via a terminal to the management network 366) or a mixture of both. The system management server 388, 399 may also be directly connected to a management terminal 365 for direct operator control of the particular system management server. In order that the above functions be
20 implemented in the system management server, the system management server may be configured to have a level of intelligence appropriate to the making of decisions and controlling of failover and deployment.

Thus there has been described an arrangement wherein management and
25 service functions are distributed in a hierarchical manner across various levels of a modular computer system. Through the implementation of a hierarchical arrangement of service functions, any given service function may be implemented at a most appropriate level. For example, a top level service controller need not be concerned with environmental monitoring within an individual processing cartridge, whereas the

same top level service controller may have responsibility for overall management for a plurality of shelves 41. As the skilled addressee will appreciate, a number of alterations and modifications may be made to the above described arrangement without departing from the scope of the present invention. Also, a number of the
5 above described components and functions may be replaced with or supplemented by equivalent components and functions without departing from the scope of the invention.

Although the embodiments above have been described in considerable detail,
10 numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such variations and modifications.

CLAIMS

1. A server system comprising a server blade carrier carrying a plurality of removably mountable server blades, each blade including a blade service controller operable to perform monitoring and management functions in the blade, the carrier comprising at least one carrier service processor operably connected to the blade service controllers of blades mounted in the carrier to provide higher level management functions and configured to communicate management functions to and from the blades.
2. The server system of claim 1, wherein the blade carrier comprises at least two carrier service processors.
3. The server system of claim 2, wherein said at least two carrier service processors are operable in a redundant manner.
4. The server system of claim 3, wherein a first service processor is operable as a master service processor and a second service processor is operable as a backup service processor.
5. The server system of any preceding claim, wherein the carrier comprises a power supply unit operably connected to supply power to the carrier service processor and to the blade servers mounted in the carrier and wherein the carrier service processor is configured to provide management functions to the power supply unit.
6. The server system of any preceding claim, wherein the carrier comprises a switch operably connected to perform distribution of information signals within the carrier and wherein the carrier service processor is configured to provide management functions to the switch.

7. The server system of claim 6, wherein the carrier service processor and the switch are arranged into a unitary combined switch and service processor module.
- 5
8. The server system of any preceding claim, wherein the blade service controller is operable to monitor an environmental variable of the blade comprising the blade service controller.
- 10 9. The server system of claim 8, wherein the environment variable includes at least one of a processor temperature, an ambient temperature, a voltage level and a fan rotation speed.
10. The server system of any preceding claim, wherein the blade service controller is operable to monitor the operation of an operating system of the blade.
- 15
11. The server system of any preceding claim, wherein the blade service controller is operable to provide an out of band interface between the blade and the carrier service processor.
- 20
12. The server system of any preceding claim, wherein the blade service controller is operable to control a processor reset function.
13. The server system of any preceding claim, wherein the carrier service processor is operable to monitor an environmental variable of the blade carrier.
- 25
14. The server system of any preceding claim, wherein the carrier service processor is operable to cause a restart of a blade in response to a blade fault.

15. The server system of any preceding claim, wherein the carrier service processor is operable to cause a blade to be removed from service in response to a blade fault.
- 5 16. The server system of any preceding claim, wherein the carrier service processor is operable to report the removal from service of the blade to a higher level management provider.
- 10 17. The server system of any preceding claim, comprising at least one blade carrier and a system management server, the system management server being operably connected to the carrier service processors of the blade servers mounted in the carrier to provide high level management functions.
- 15 18. The server system of claim 17, comprising at least two blade carriers.
19. The server system of claim 17 or claim 18, comprising at least two system management servers.
- 20 20. The server system of claim 19, wherein said at least two system management servers are operable in a redundant manner
21. The server system of claim 19, wherein said at least two system management servers are operable as a clustered group.
- 25 22. The server system of any of claims 17 to 21, comprising a management network, the management network being operably connected to the system management server to provide high level management functions.

23. The server system of any of claims 17 to 22, comprising a management terminal, the management terminal being operably connected to the system management server to provide high level management functions.
- 5 24. The server system of any of claims 17 to 23, wherein the system management server is operable to perform deployment of software to a blade.
25. The server system of any of claims 17 to 24, wherein the system management server is operable to perform hardware fault failover of blades.
- 10 26. The server system of any preceding claim, wherein the blades are configured to be hot-swappable.
27. The server system of any preceding claim, wherein the blades are configured
15 as field replaceable units.
28. The server system of any preceding claim, wherein the carrier service processor is configured as a field replaceable unit.
- 20 29. The server system of any preceding claim, wherein the blade carrier is a rack-mountable shelf.
30. A method of providing management functions within a server system having a server blade carrier carrying a plurality of removably mountable server blades,
25 the method comprising:
performing monitoring and management functions at a blade service controller within each blade;
performing higher level management functions at at least one carrier service processor within the carrier;

providing a communications channel connecting the carrier service processor to each blade service controller; and
communicating management functions between the carrier service processor and the blade service controllers.

5

31. The method of claim 30, further comprising:

performing higher level management functions at a system management server remote the carrier;

providing a communications channel connecting the system management

10

server to the carrier service processor; and

communicating management functions between the system management server and the carrier service processor.

1/25

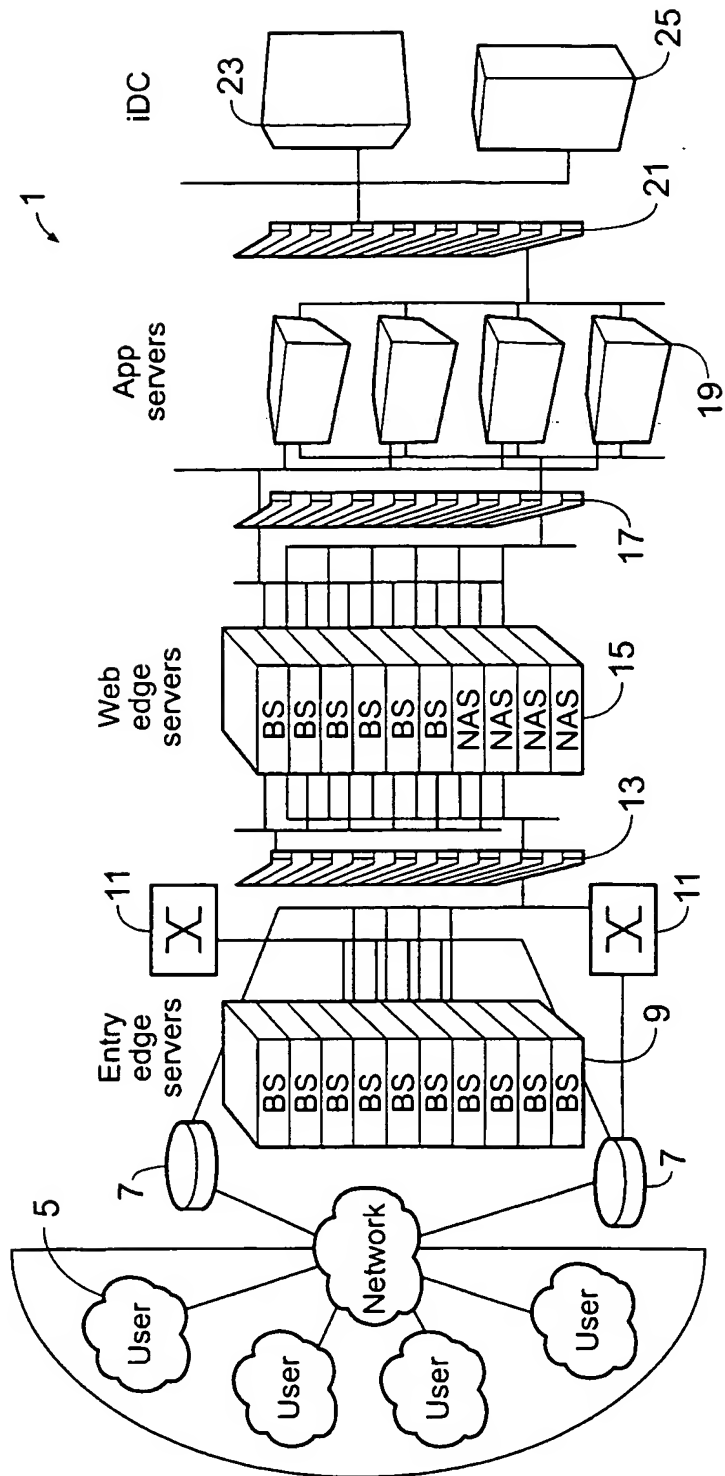


Fig. 1

2/25

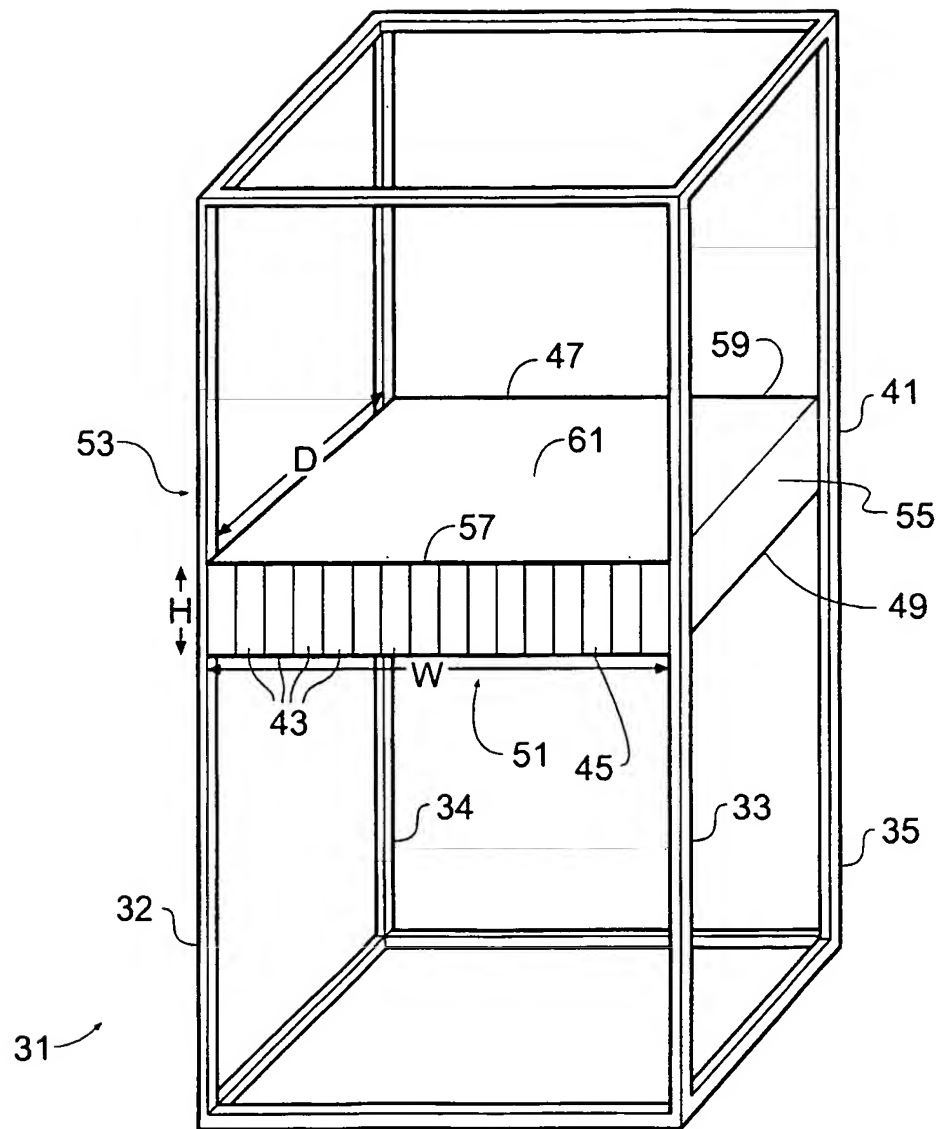


Fig. 2

3/25

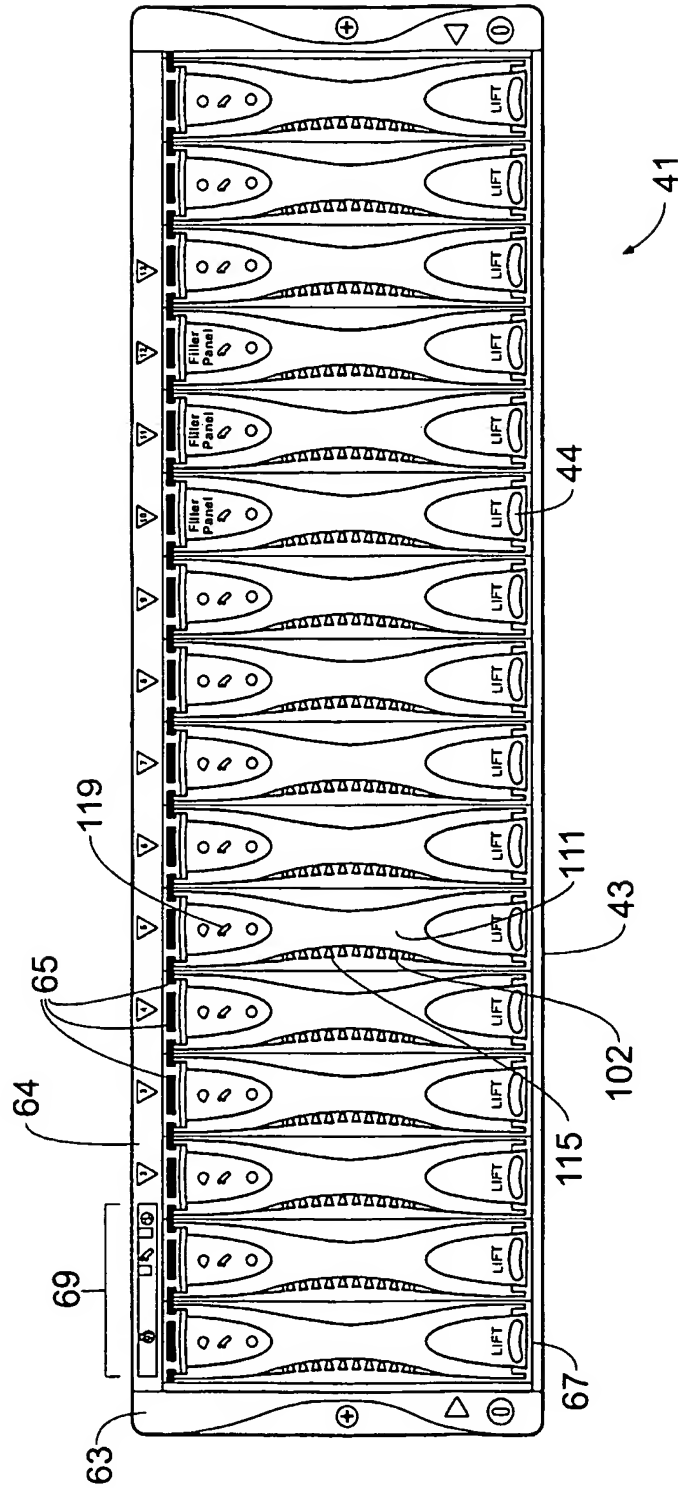


Fig. 3

4/25

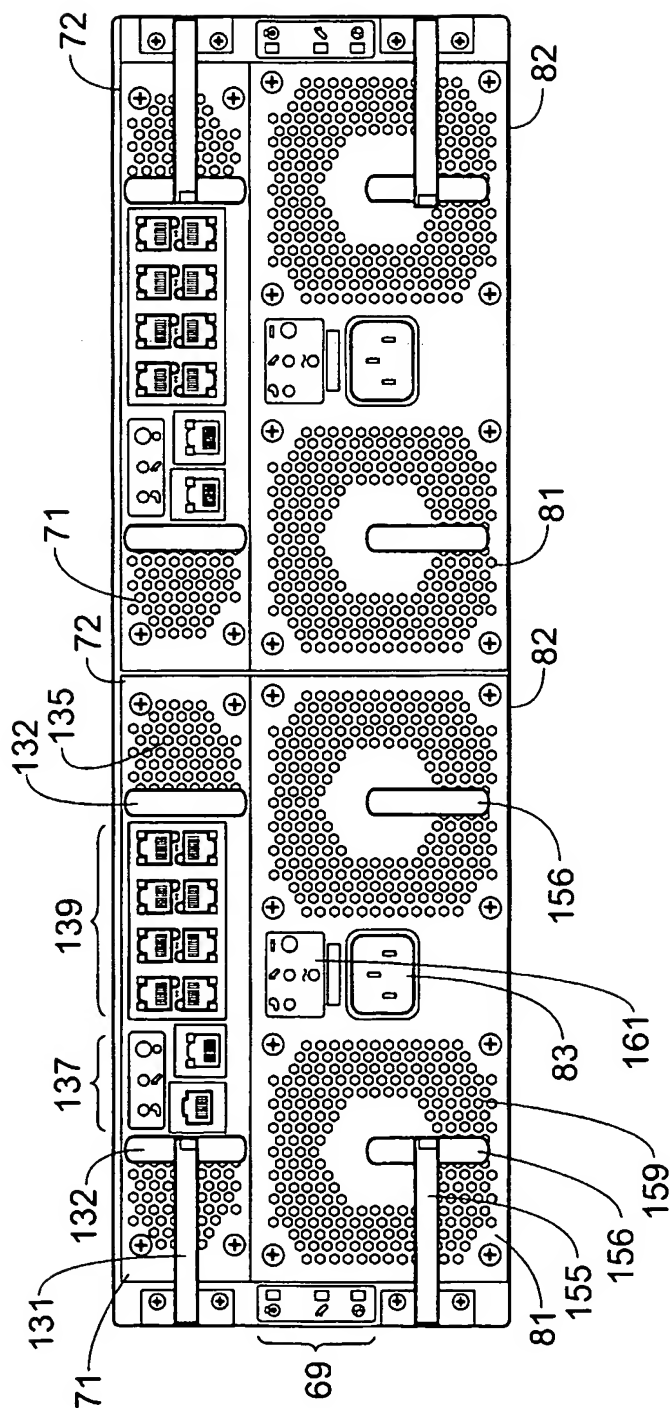


Fig. 4

5/25

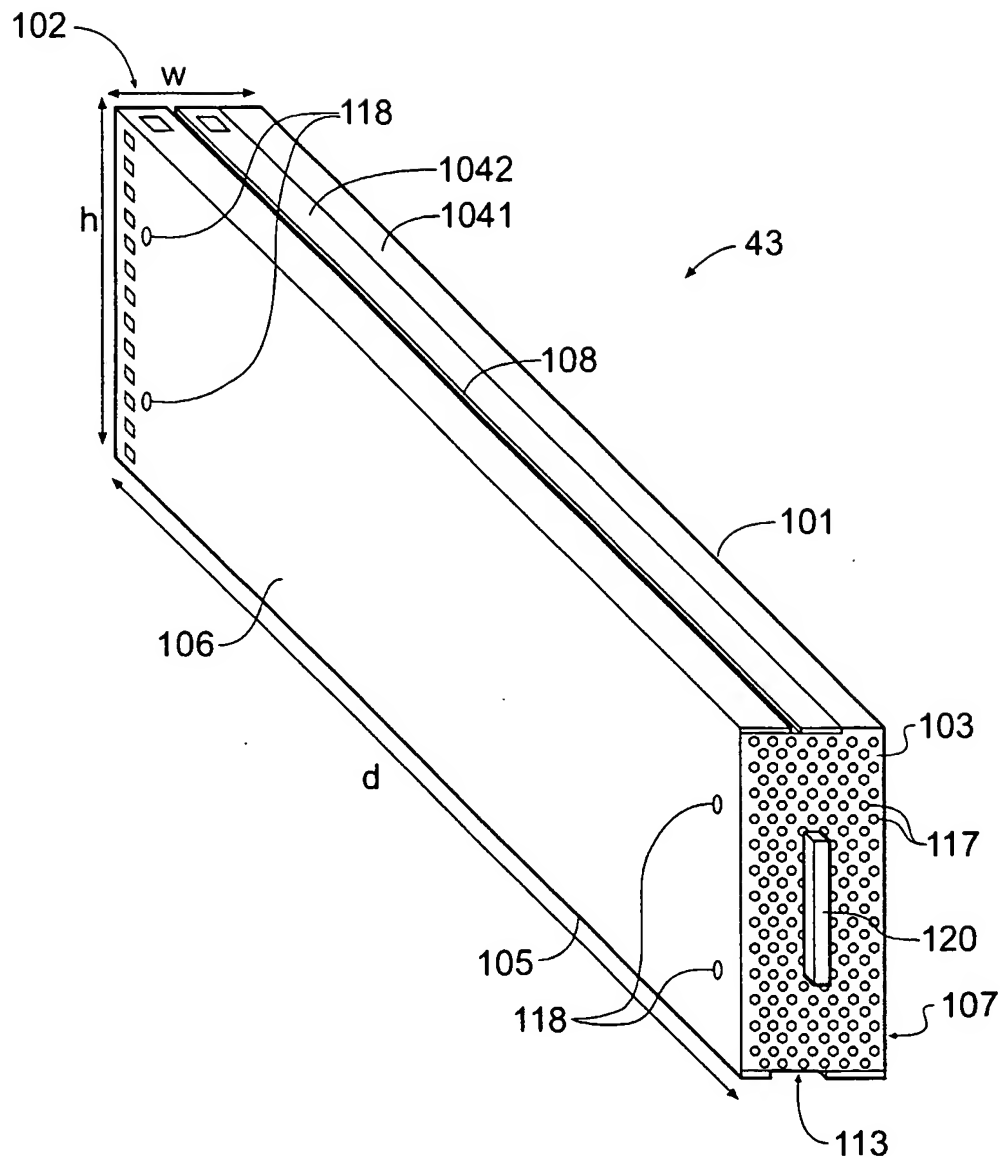


Fig. 5A

7/25

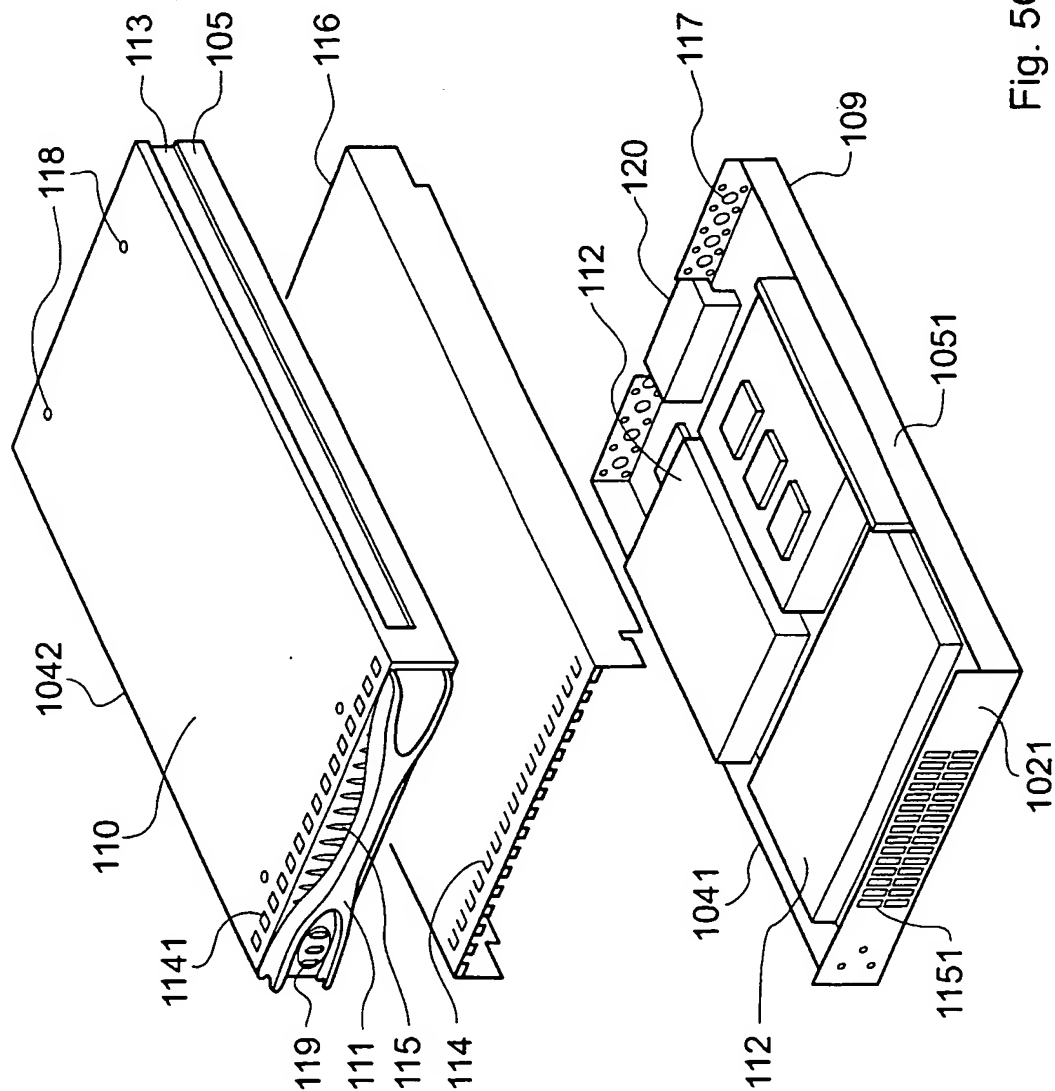


Fig. 5C

9/25

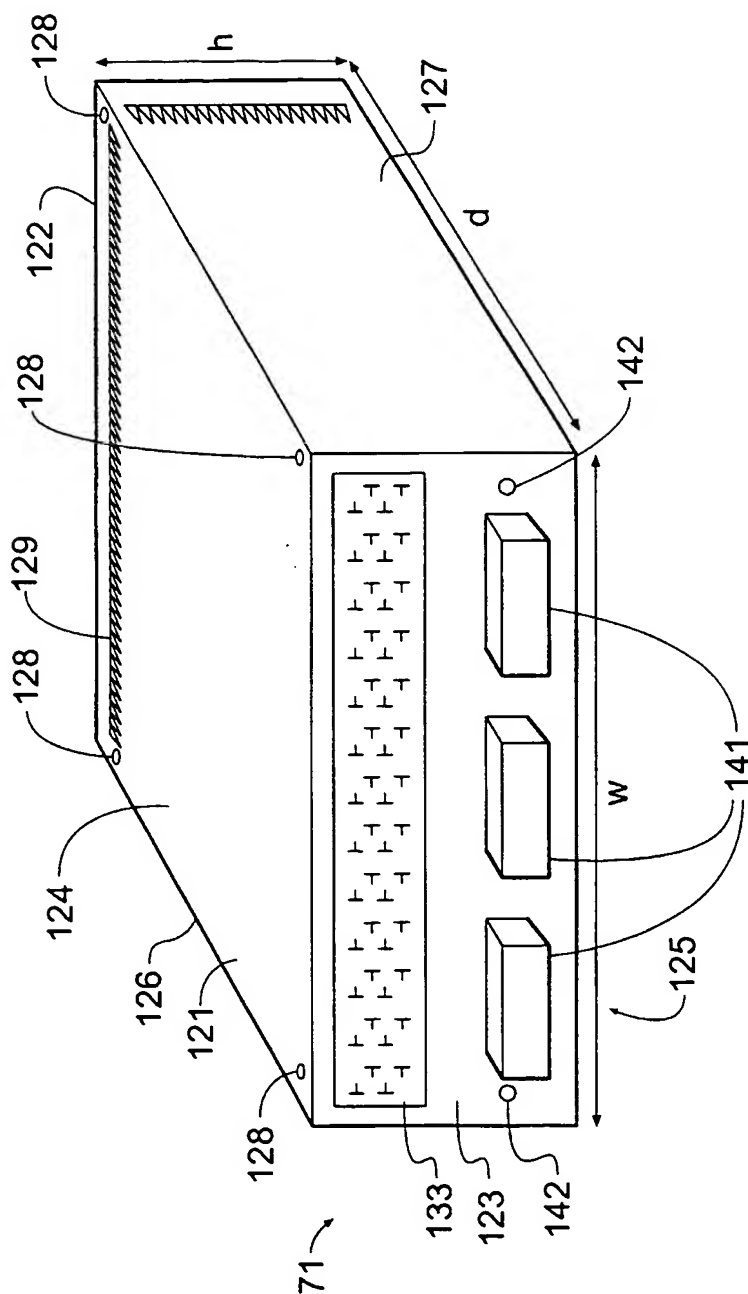


Fig. 6

10/25

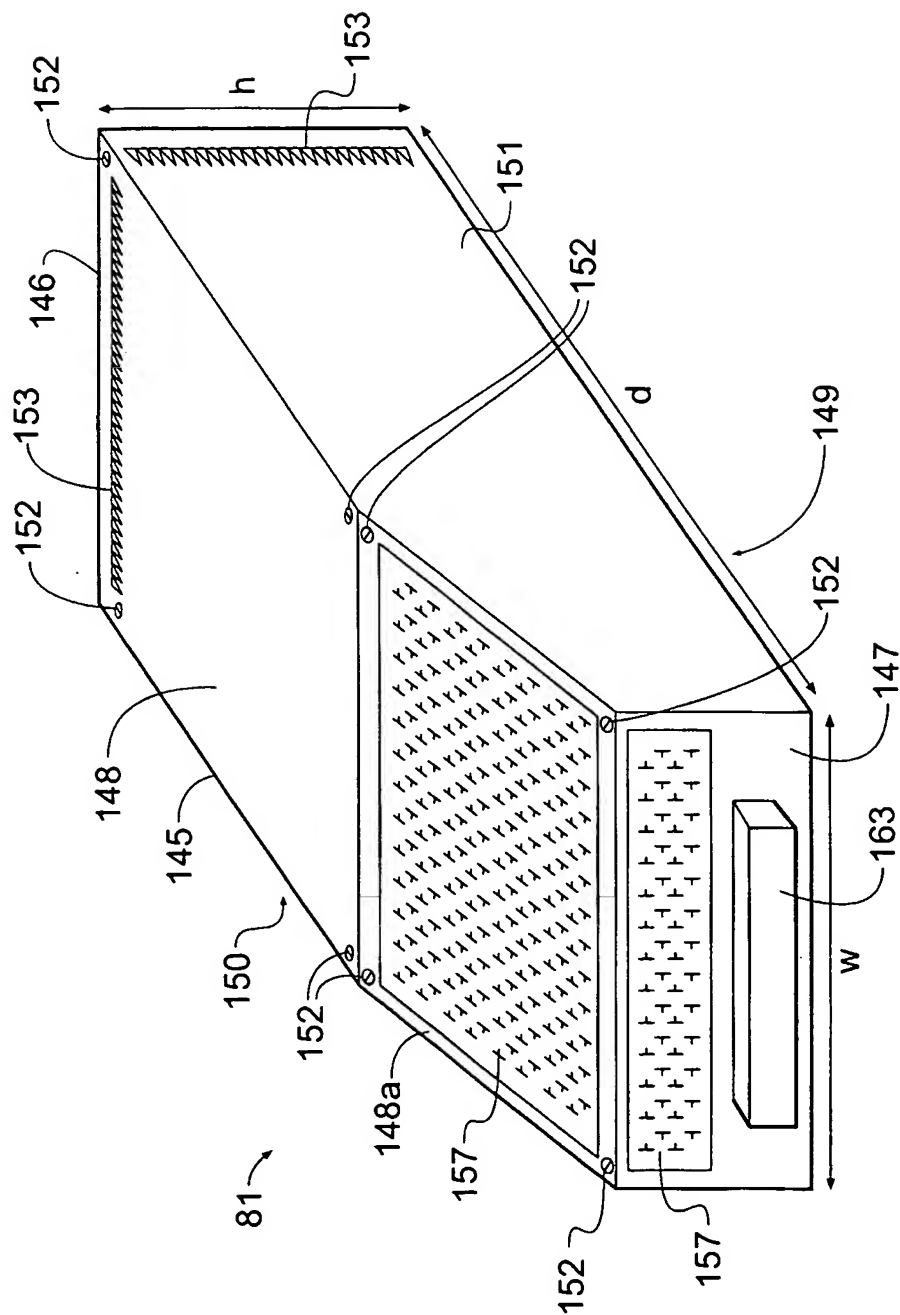


Fig. 7

11/25

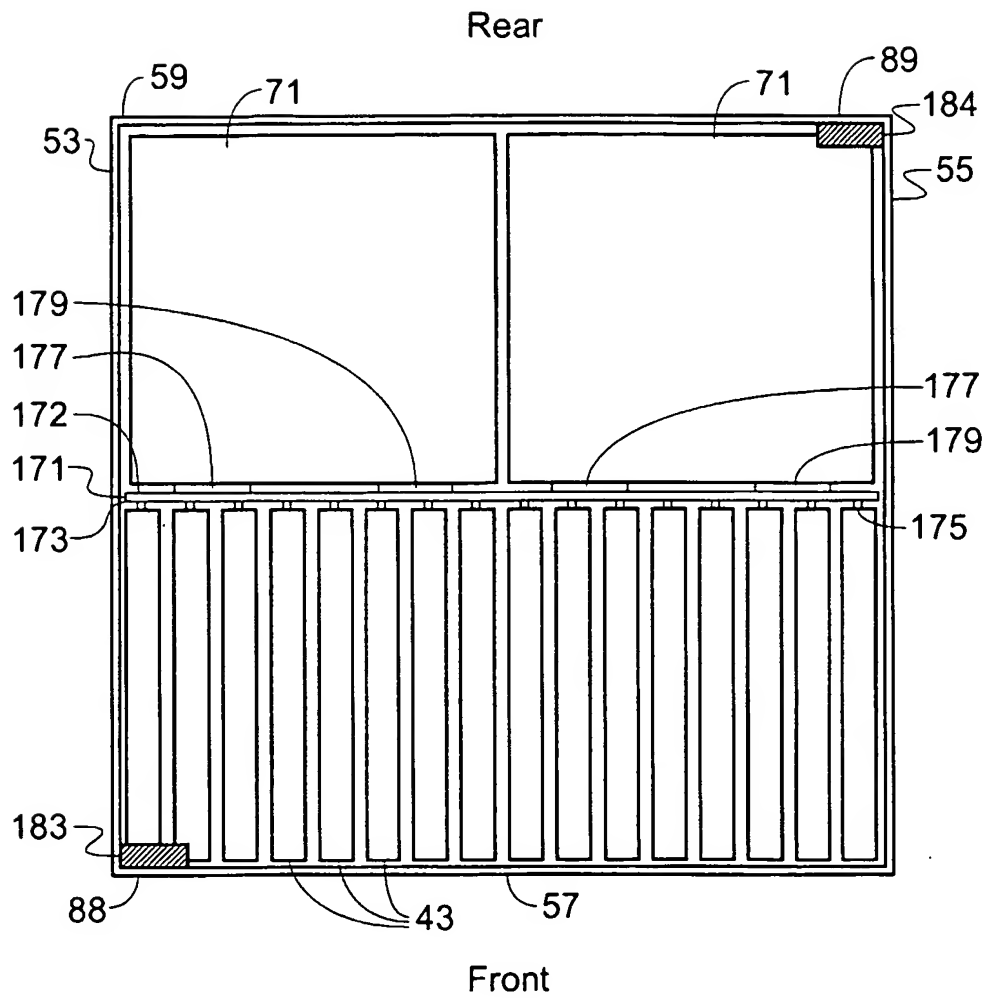


Fig. 8A

12/25

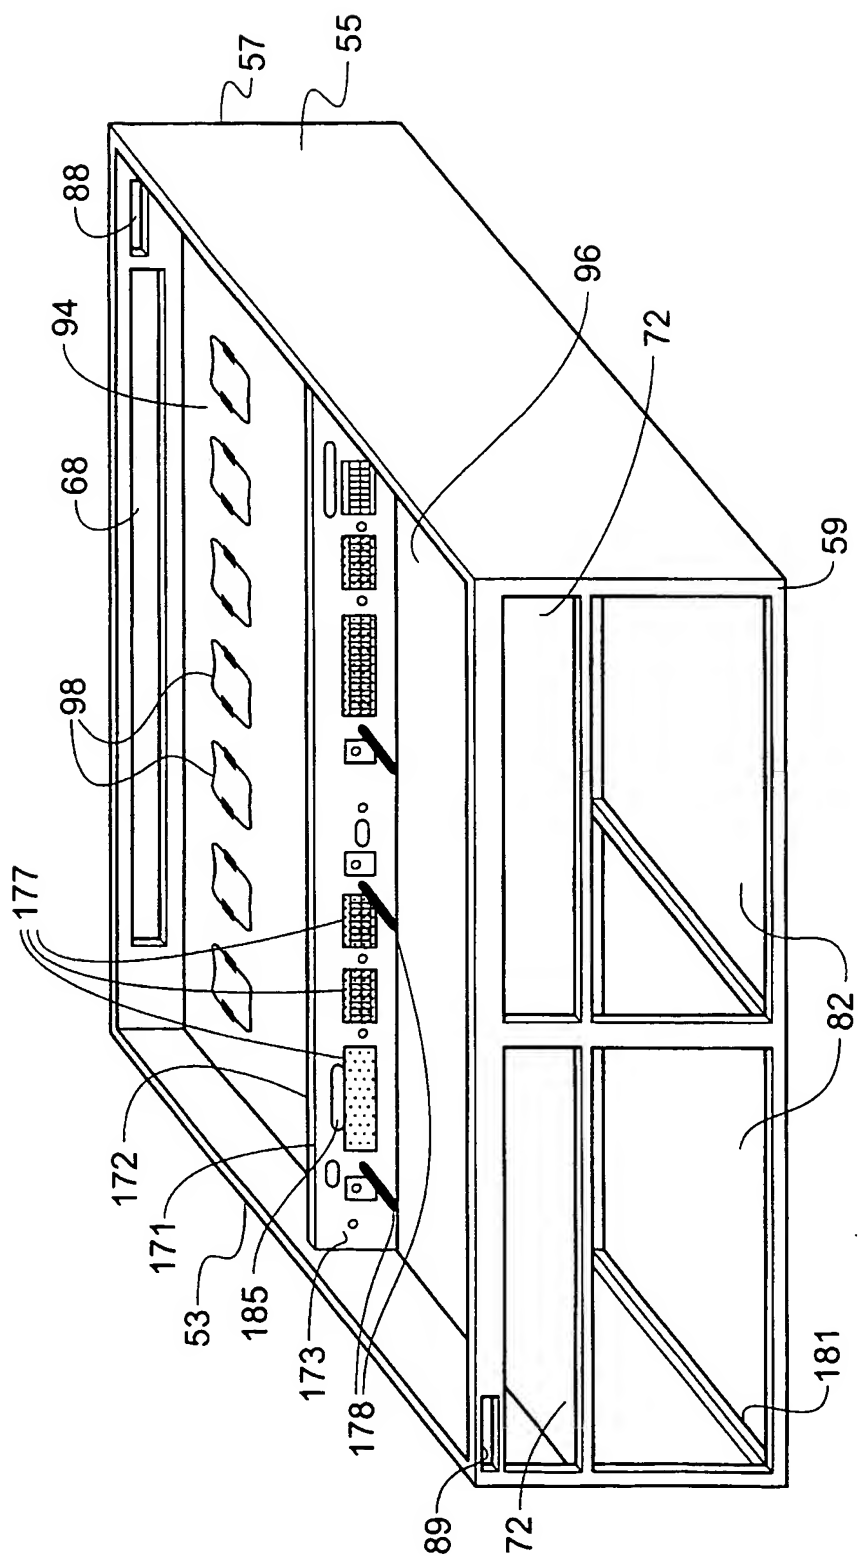


Fig. 8B

13/25

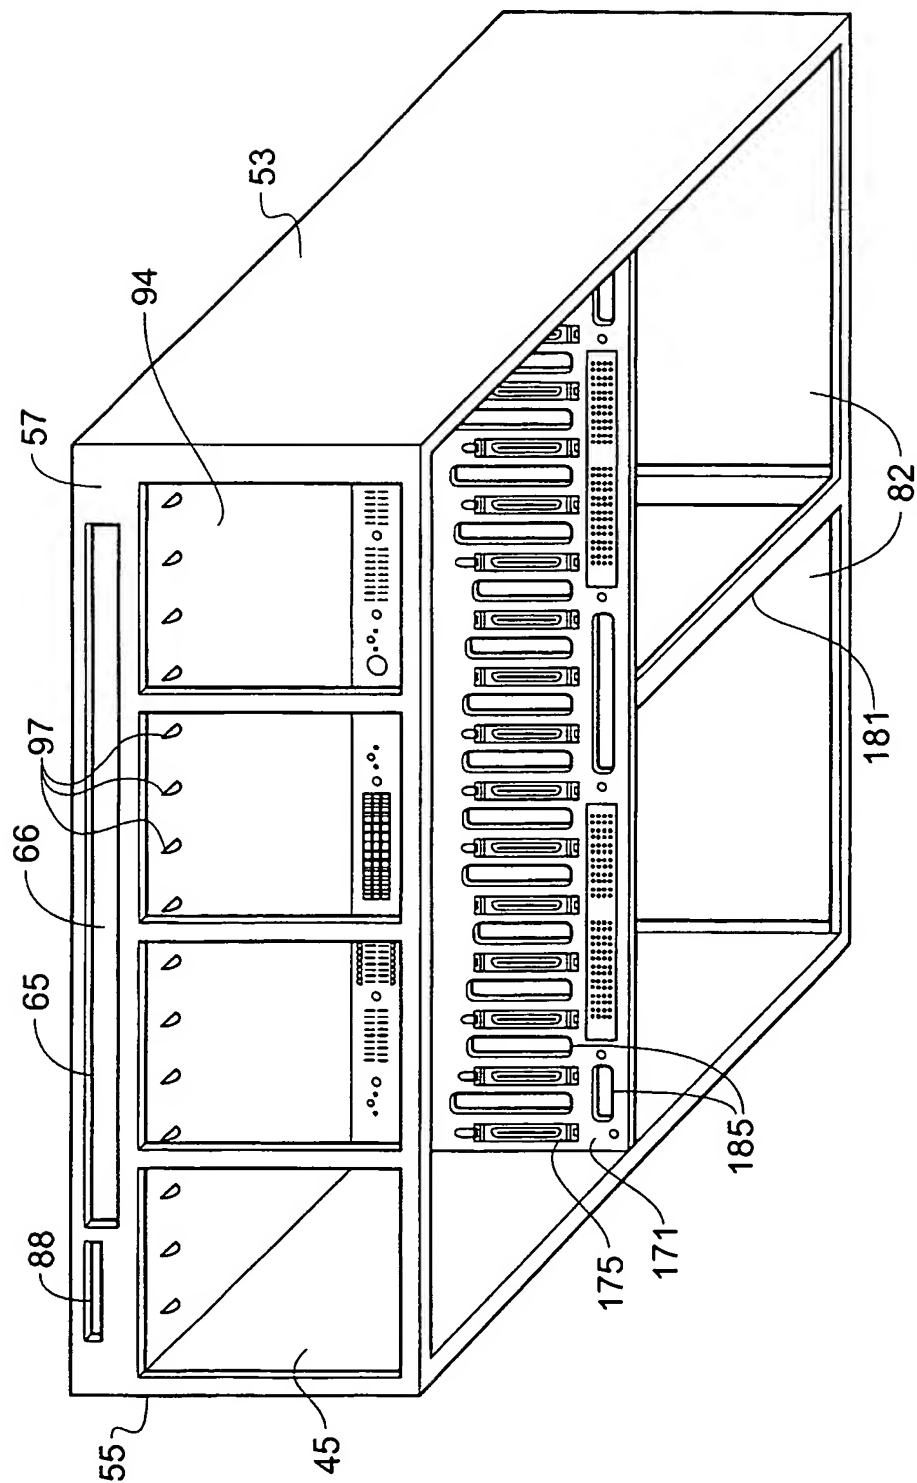


Fig. 8C

14/25

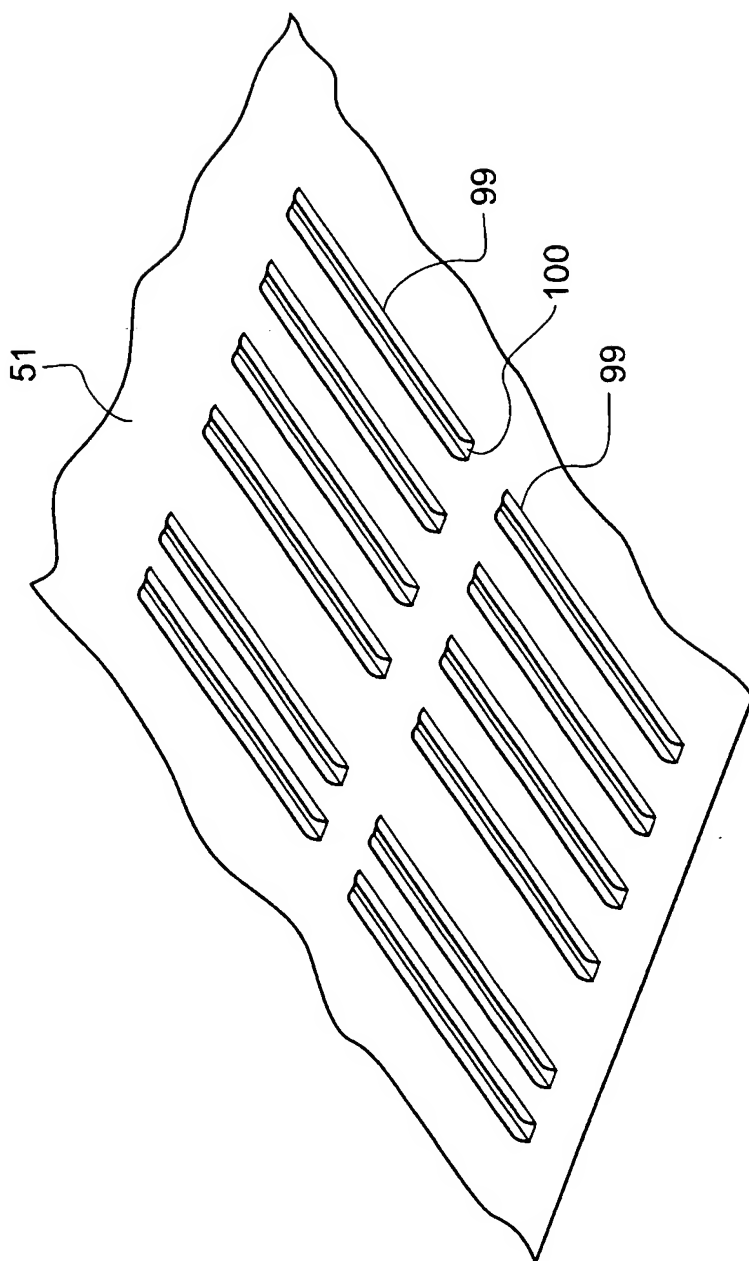


Fig. 8D

15/25

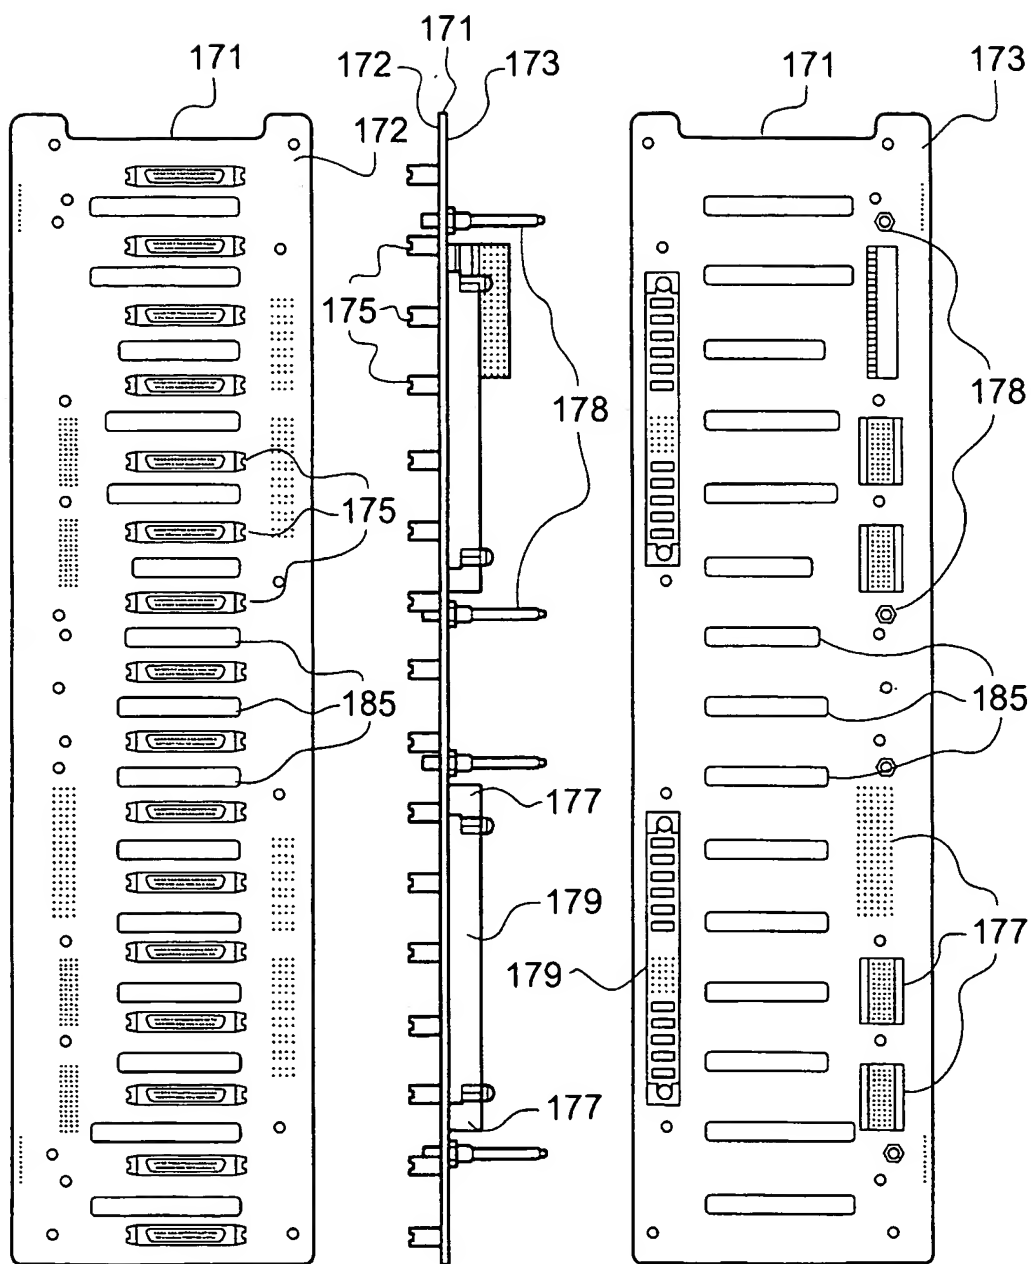
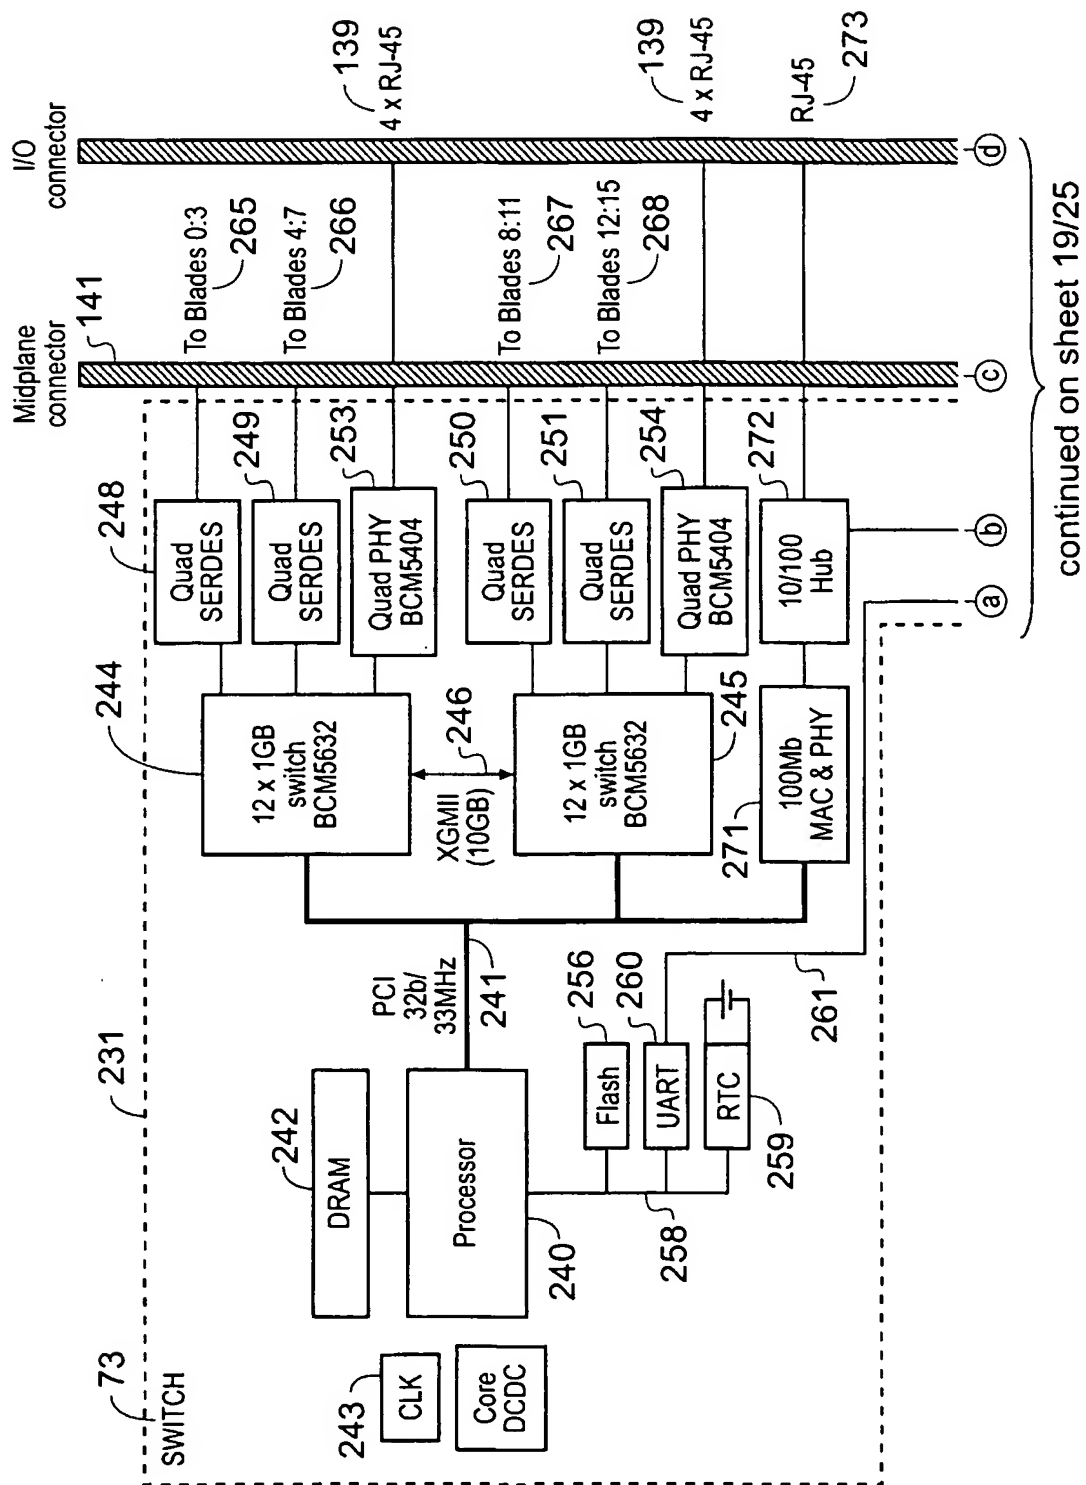


Fig. 9A

Fig. 9B

Fig. 9C

18/25



19/25

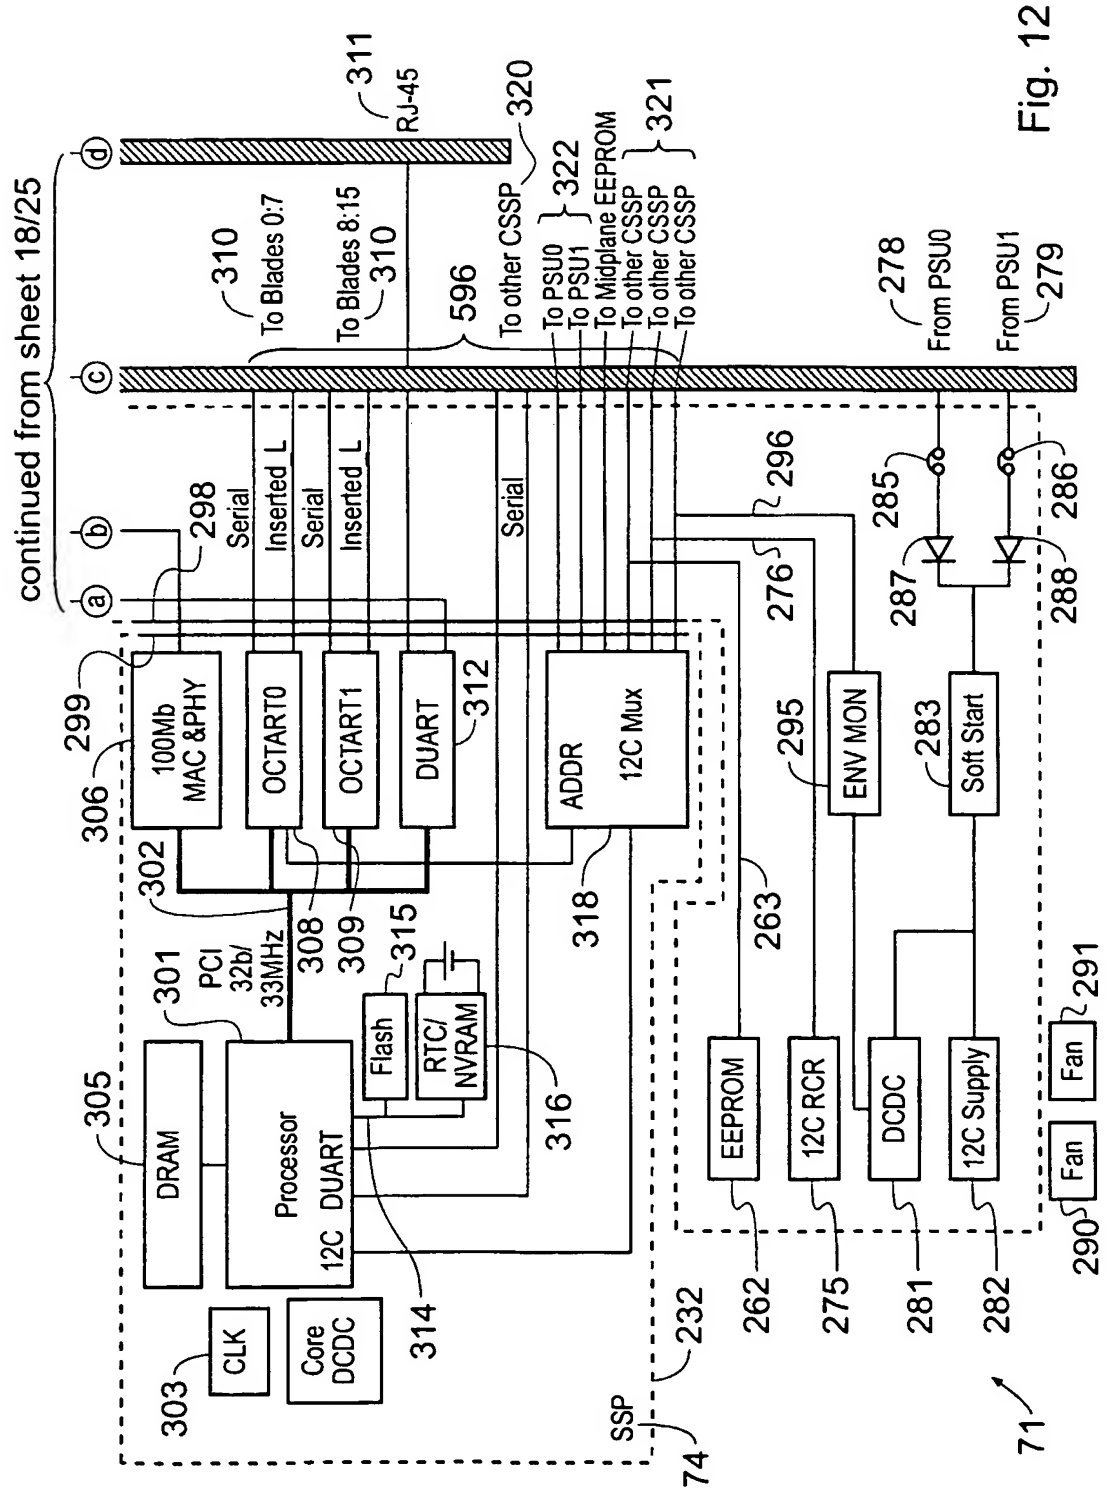


Fig. 12

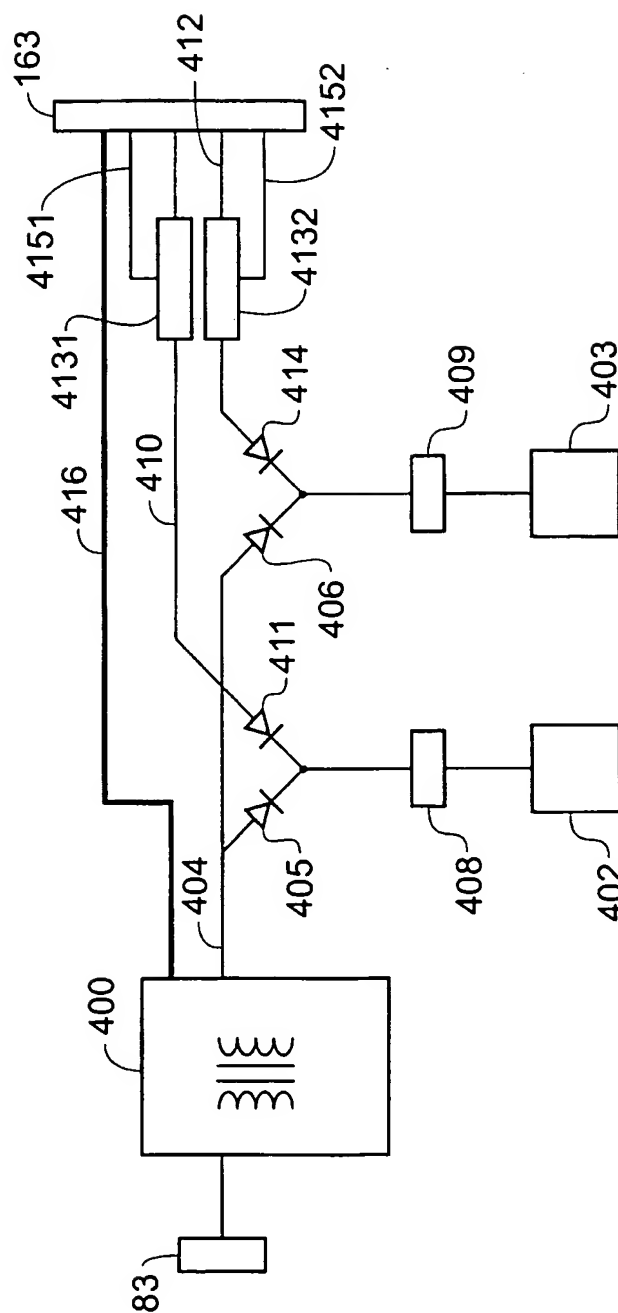


Fig. 13

21/25

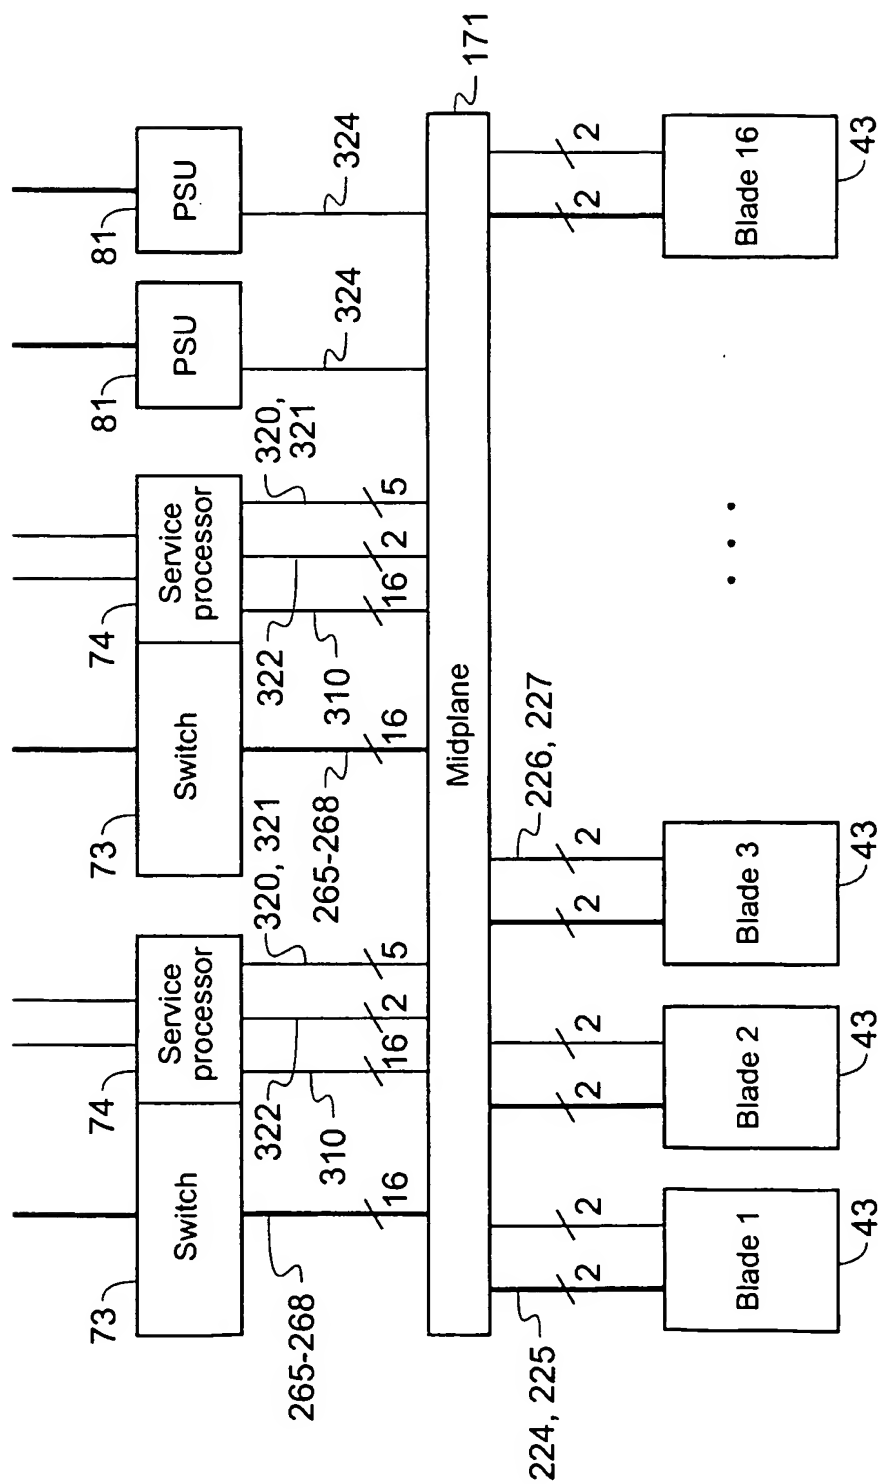


Fig. 14

22/25

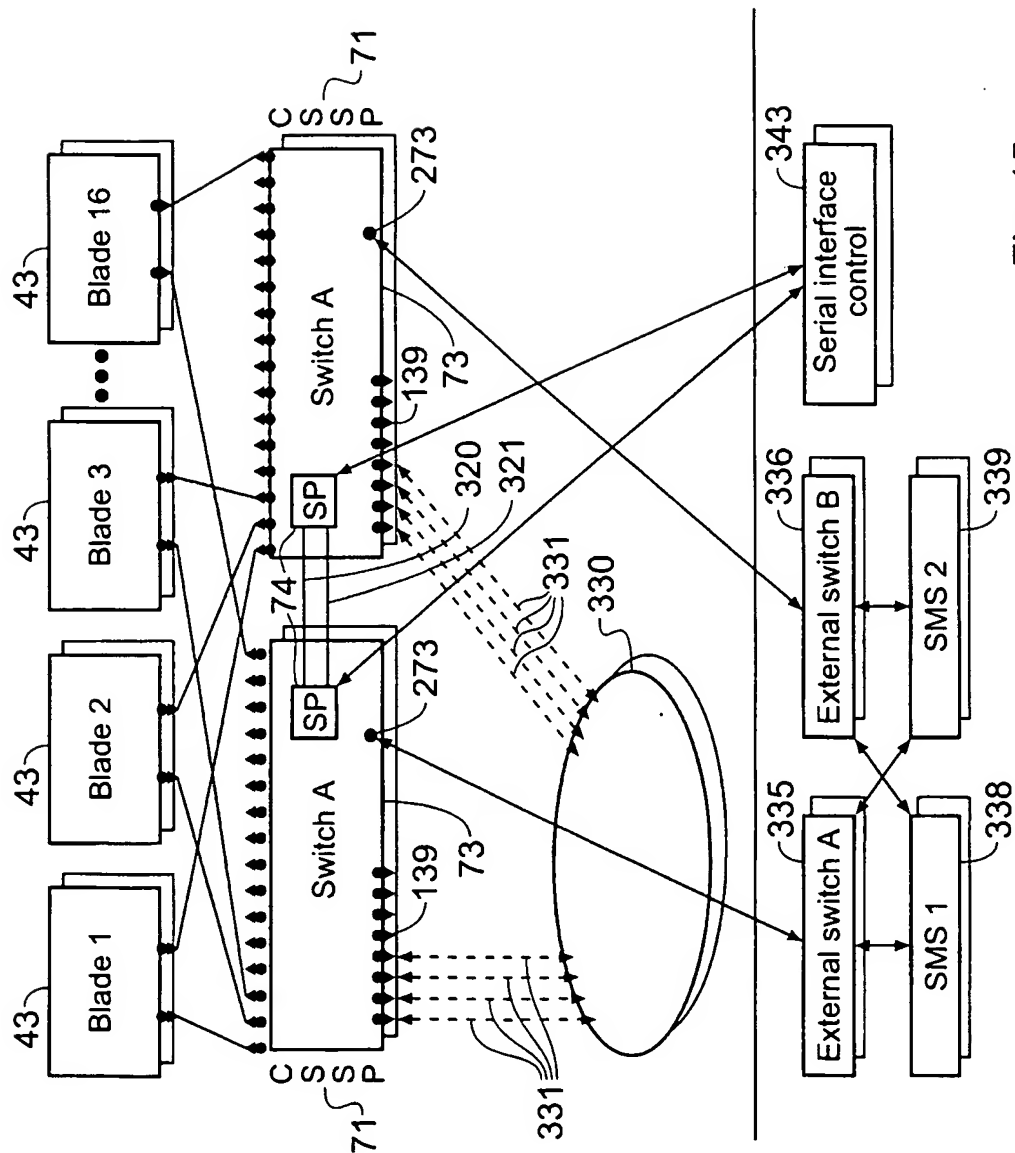


Fig. 15

23/25

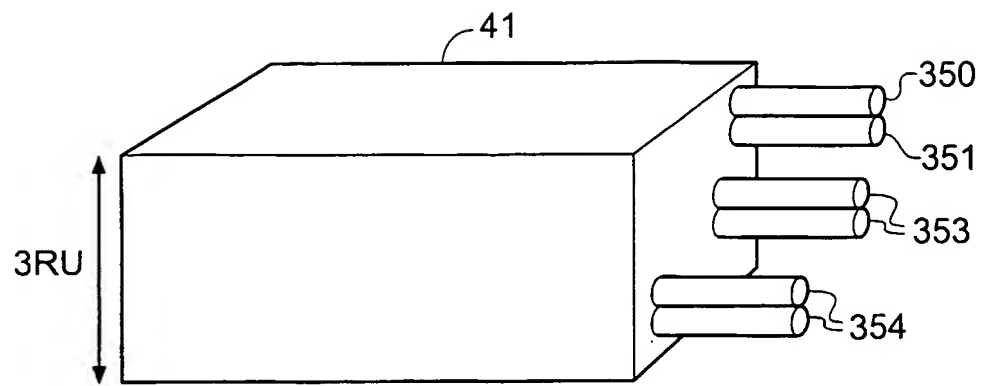


Fig. 16

24/25

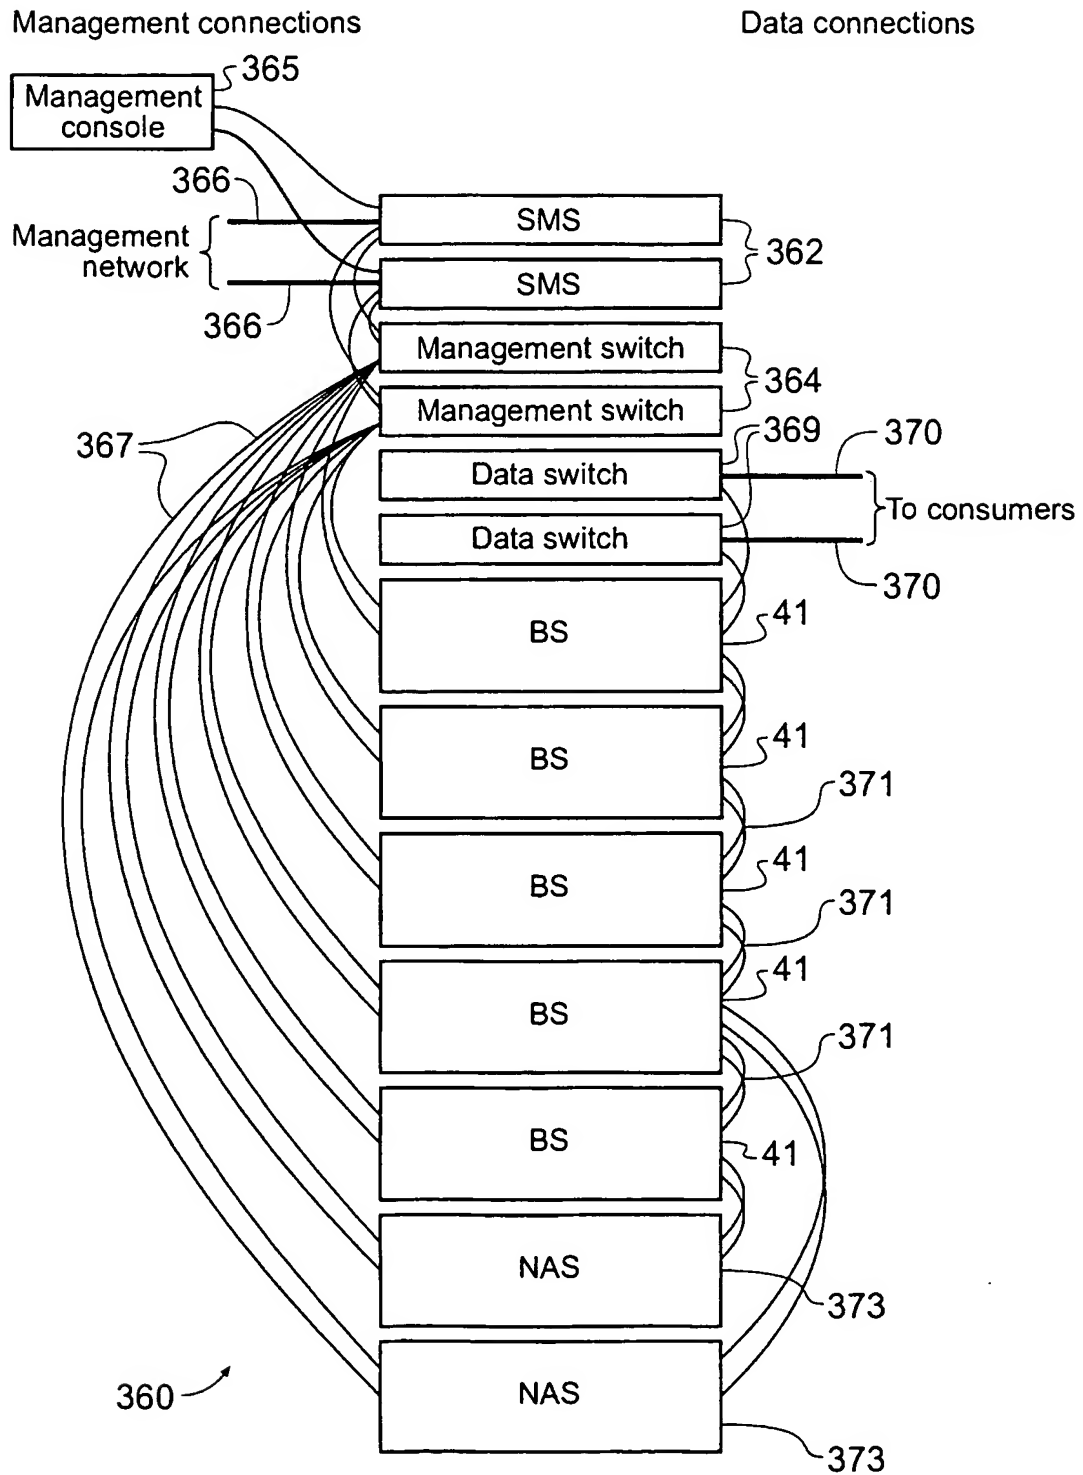


Fig. 17

25/25

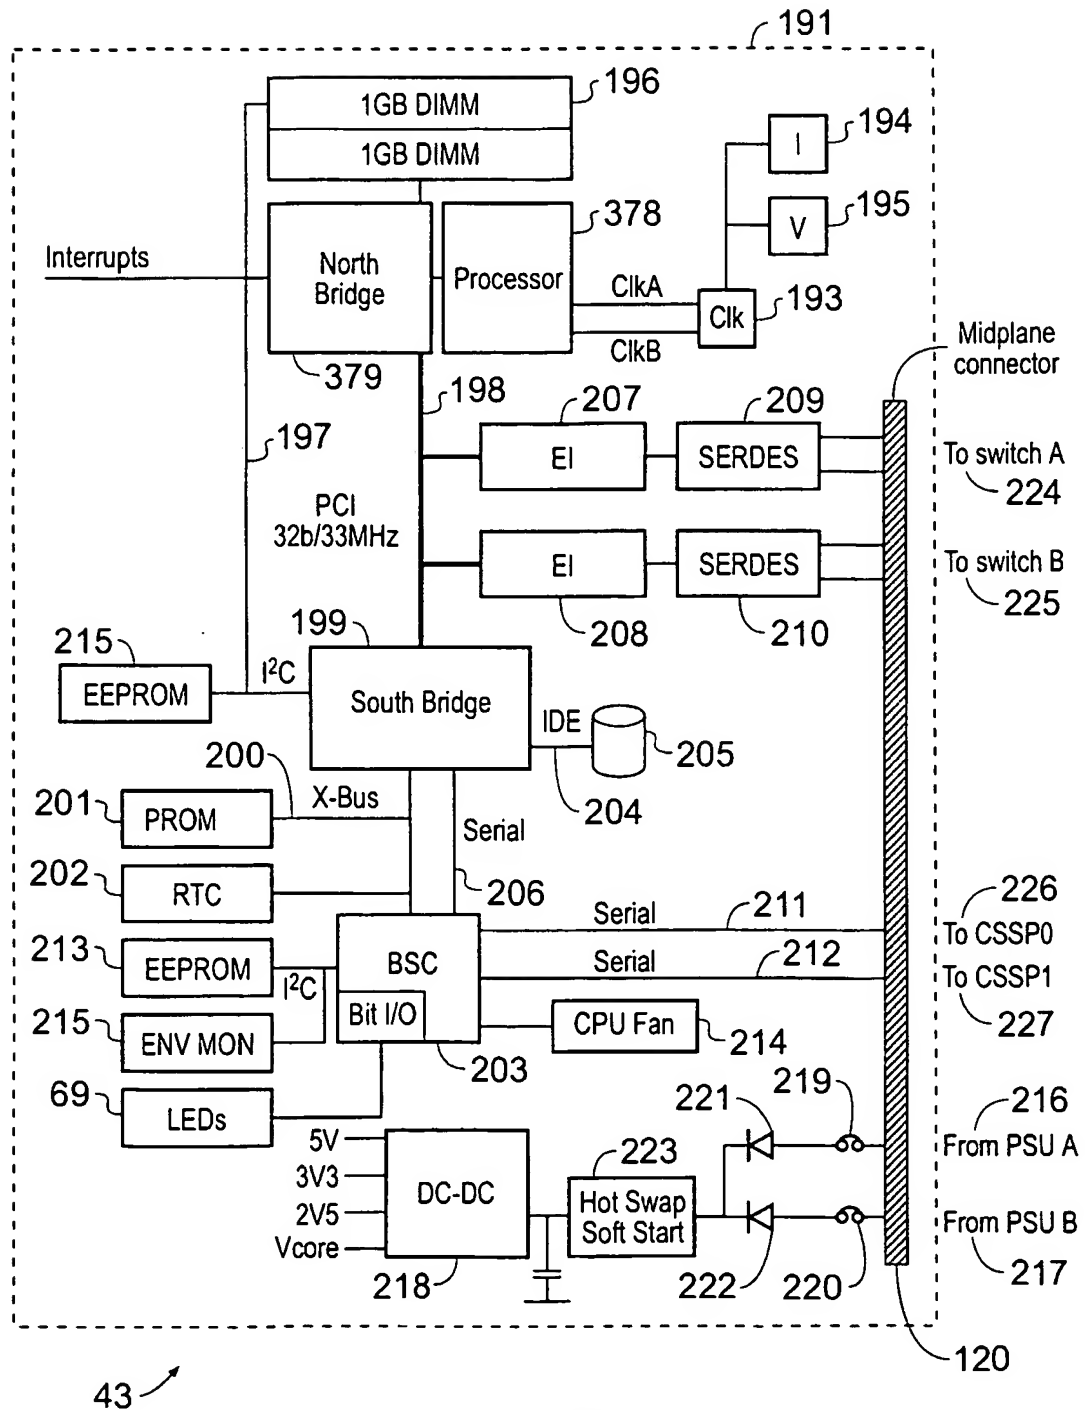


Fig. 18